

# Integration of Sharemind MPC into Carbyne Stack

Technical Report

Version 1.0

31.01.2023

D-2-502

Public

Copyright ©2023

Cybernetica AS

All rights reserved. The reproduction of all or part of this work is permitted for educational or research use on condition that this copyright notice is included in any copy.

Cybernetica research reports are available online at

<https://research.cyber.ee/>

Mailing address:

Cybernetica AS

Mäealuse 2/1

12618 Tallinn

Estonia

# Table of Contents

---

<b>1 Introduction</b>	<b>4</b>
1.1 Purpose	4
1.2 Scope	4
1.3 Overview	4
1.4 Definitions, Acronyms and Abbreviations	4
<b>2 Comparison of the Two Platforms</b>	<b>6</b>
2.1 Services	6
2.2 Environment	7
2.3 Deployment and Setup	7
2.4 Client Interface	8
2.4.1 I/O	8
2.4.2 Session Management	8
2.4.3 Access and Policy Control	8
<b>3 Motivation</b>	<b>9</b>
<b>4 Integration Plan</b>	<b>10</b>
4.1 Goals	10
4.2 High Level Vision	10
4.3 Development Roadmap	11
4.3.1 Milestone 1: A Minimal Runtime	11
4.3.2 Milestone 2: Task Arguments and Published Values	13
4.3.3 Milestone 3: Distributed Tabular Databases	13
<b>Appendix A Extra Milestones</b>	<b>16</b>
A.1 Task Scheduling	16
A.2 Transactional Distributed Databases	16
A.3 Gateway Applications	16

# 1 Introduction

---

## 1.1 Purpose

For a year leading up to this report has the Carbyne Stack open source project had the attention of the Sharemind MPC team. The project's significance to us is multifaceted. First, it recognizes the challenges of MPC deployments and offers a cloud-native stack that is easy to deploy and configure. Secondly, it proposes an API which is not necessarily specific to the underlying MPC backend. Finally, its potential for scaling MPC over a cluster. These qualities justify adapting the core tools and protocols of Sharemind MPC to run inside Carbyne Stack's managed environments. This document seeks to answer the *what, why, and how* in regard to this integration.

The first part of this report expands on the publicly available documentation of Carbyne Stack based on insight from various secondary resources (source code, meetings, etc.). It does so in a comparative manner, describing how certain details compare to Sharemind MPC, with the goal of making it more comprehensive to ones who are more familiar with one of the platforms, but less so with the other. The second part outlines the motivations behind the integration. The final part explores technicalities of a possible integration, the purpose of which is to lay groundwork for the starting of development and invoke discussions.

A reader of this report is expected to have a general understanding of secure multi-party computation and core concepts of both the Carbyne Stack<sup>1</sup> and Sharemind MPC<sup>2</sup> platforms.

## 1.2 Scope

This report first and foremost explains the plan of integrating the two MPC platforms, from conception to a state of supporting the main features of Sharemind MPC. Additionally, some extended goals are outlined.

Details which are deemed significant for the integration are explored on a more technical level, while some questions are left open-ended. This report is subject to new revisions as the Carbyne Stack project matures over time, filling in gaps that can not be planned for at this time.

## 1.3 Overview

Chapter 2 introduces and compares the platforms. Chapter 3 explains the integration's motivation. The draft plan along with technical aspects and open problems is described in chapter 4.

## 1.4 Definitions, Acronyms and Abbreviations

### VCP / computing party

*Virtual Cloud Provider* – a single MPC party which provides an MPC computation service

### VC / clique

*Virtual Cloud* – a set of VCPs which jointly engage in MPC computation

### Pod

smallest deployable unit in Kubernetes; usually a single container

---

<sup>1</sup>Carbyne Stack introductory video: <https://carbynestack.io/getting-started/>

<sup>2</sup>Overview of Sharemind MPC: <https://docs.sharemind.cyber.ee/2022.03/prologue>

**(MPC) task**

execution of a single MPC program

**ACL**

Access-control list

**IaC**

Infrastructure as Code

**FaaS**

Function-as-a-Service or serverless computing

**ACID**

Atomicity, Consistency, Isolation, Durability – defining properties of a database transaction

## 2 Comparison of the Two Platforms

---

Carbyne Stack is a software stack intended to fill the market void of an enterprise focused MPC platform that is open-source and community maintained [1]. It can be characterized as a set of middleware and supporting facilities for clients to invoke MPC programs and store encrypted data over multiple Carbyne Stack deployments acting as computing parties. It is noteworthy that

1. it follows a client-server model of MPC, in which clients – data providers and data recipients – do not partake in the computation, but rather offload the computation along with any encrypted inputs to the computing parties' clusters;
2. the platform does not itself supply, and is not technically constrained to any MPC backend – presently it leverages the MP-SPDZ [2] project for its protocols and runtime.

At a glance, Sharemind MPC bears resemblance to Carbyne Stack with regard to the client-server model, its general purpose and functionality. It does however incorporate its own runtime and set of MPC protocols, which are tightly coupled to the platform as a whole.

With both platforms aiming to provide a similar MPC pipeline, they share a number of architectural challenges. The following section compares the differences in solving some of these challenges based on the implementations at the time of writing.

### 2.1 Services

Carbyne Stack adopts a *cloud native*<sup>1</sup> approach which constitutes a microservices architecture. As with many other cloud native projects its services largely rely on the Kubernetes API and other technologies of the CNCF<sup>2</sup> ecosystem. It is split into the following services:

1. Ephemeral is a FaaS that executes MPC computation processes per clients' requests. Knative Serving<sup>3</sup>, a solution for FaaS within Kubernetes, is used to scale Ephemeral instances down to zero when idle, and transparently spawn pods as requests arrive. This effectively distributes MPC computation throughout the cluster.
2. Network Controller lives in Kubernetes' control plane, monitoring the cluster for new MPC tasks, for which it exposes a cluster port. Istio's Virtual Services<sup>4</sup> are used to create and manage those routes.
3. Discovery is a stateful service that is used by Ephemeral instances to register the cluster endpoint assigned to the task and fetch the public endpoints of the neighbouring clusters' task. The Discovery service of each party communicates with its counterparts to learn the TCP ports for the tasks.
4. Amphora is the service for storing and receiving secret shared values from cloud or in-cluster object storage.
5. Castor is the storage service for correlated randomness.
6. Klyshko is responsible for running the *offline phase* of MPC to generate correlated randomness.

---

<sup>1</sup>The concept of cloud native: <https://aws.amazon.com/what-is/cloud-native/>

<sup>2</sup>CNCF - *Cloud Native Computing Foundation*: <https://www.cncf.io/>

<sup>3</sup>Knative Serving: <https://knative.dev/docs/serving/>

<sup>4</sup>Virtual service: <https://istio.io/latest/docs/concepts/traffic-management/#virtual-services>

Neither a microservices architecture or distribution of computation over nodes has been embraced by Sharemind MPC. Instead the central server-side component, called the application server, has a monolithic architecture. It does however host a loadable module system to add new functionality to the virtual machine running the MPC program. Some modules act as services within the application server which can be called from the MPC program's bytecode and other modules.

Similarly to Ephemeral, the application server on its own is only responsible for handling client sessions, starting virtual machines, and managing the bytecode programs. Facilities such as Carbyne Stack's Network Controller and Discovery are not necessary due to a fixed network topology that is coordinated and configured ahead of time – only one port is utilized, with messages routed to respective sessions and facilities within the application server, referred to as multiplexing.

`tabledb` and `keydb` modules provide MPC programs access to tabular and key-value databases. These modules have a database management layer which synchronizes write operations with other computing parties. Amphora's key-value nature expects keys as UUIDs supplied by the client and therefore does not require synchronization of writes.

Castor and Klyshko have no counterparts in Sharemind MPC for the time being. Future development in active security protocols may alter the fact.

## 2.2 Environment

Either platform has distinct requirements for its environments. With the cloud native approach, Carbyne Stack relies on Kubernetes to provide a well defined environment. It follows that the platform comprises a set of declarative configurations for Kubernetes, describing among others the storage and networking requirements for the services. A Kubernetes cluster that can provide a public IP and persistent storage should therefore be sufficient to host the deployment.

The Sharemind MPC application server and accompanying libraries are generally installed as packages onto supported Linux distributions running on bare metal or a virtual machine. Some dependencies rely on packages available to the host operating system through its official repositories, although containerization or virtualization help avoid this.

While Kubernetes clusters are often already configured for external access, the Sharemind MPC application server requires a public IP or its main TCP port to be forwarded to the host from outside the local network.

Sharemind MPC can store intermediate results in either the filesystem or a Redis datastore. The filesystem option requires the host to have a sufficient amount of reliable disk space, while the latter expects a separate Redis instance accessible from the host.

## 2.3 Deployment and Setup

Either platform requires involved cooperation and configuration for setup by the corresponding administrators. Mutually, they need to exchange the public IP their services are accessible from. With Sharemind MPC, its further necessary to supply all parties with each other's certificates and the certificates of clients, which is the basis of mutual authentication and TLS going forward. Computing parties are responsible for auditing and submitting code into their Sharemind MPC instance and coordinate the ACL configuration to be equal between one another. Carbyne Stack's authentication of clients and neighbouring parties, as well as code auditing is

yet a work in progress.

Even with a microservices architecture Carbyne Stack offers a relatively simple overall deployment with the Helm package manager, with cloud specific IaC in the works to improve the process further.

## 2.4 Client Interface

Every client accessible service in Carbyne Stack is exposed as a HTTP, with Java client libraries and CLI programs provided. Sharemind MPC utilizes a proprietary protocol over TCP for client communication – the native C++ API, CLI programs, and HTTP gateway libraries are among the available options for clients to communicate with the VC.

### 2.4.1 I/O

All input and output in Sharemind MPC is handled via arguments and published values<sup>5</sup> of tasks. By that, clients have no direct access to the persistence facilities. Carbyne Stack operates in the opposite manner – clients can only use data available to the Amphora storage in a task, with task results subsequently stored in the database. If clients upload values to Amphora or their task finishes with a result, they are left with a key corresponding to that value, which can be used to fetch the value or be used as input to a new task.

Amphora's current implementation is specific to the SPDZ protocol, that requires message authentication codes (MAC) accompany each secret share for active security. Its API facilitates the verification of sharings on fetching secrets from the VC and creation of authenticated shares on upload. Sharemind MPC does not have facilities for operations on authenticated secret shares.

As of now, the I/O possibilities of Carbyne Stack are restrictive

1. Amphora only supports secret shared integers;
2. Ephemeral can only write a single value to Amphora per task;
3. MP-SPDZ tasks have no access to data other than arguments (individual values) specified by the client, making the processing of larger datasets cumbersome.

Tasks in Sharemind MPC can read and write to tabular (HDF5) or Redis key-value databases during execution. Integers, floating point numbers, booleans, vectors and strings are supported throughout.

### 2.4.2 Session Management

In both cases, the client's task initiation triggers a long-running request, either in the form of a persistent TCP connection for Sharemind MPC or HTTP for Carbyne stack, for the duration of the task. Subsequently, the response holds the result of the computation.

### 2.4.3 Access and Policy Control

Sharemind MPC's client authorization works on both database (table/key, read/write) and program level, specifying the programs each client is allowed to execute. This enables specification of intricate MPC applications with clients of various roles.

---

<sup>5</sup>private values revealed to the client



## 3 Motivation

---

There are several justifications to undertake an integration of the Sharemind MPC runtime into Carbyne Stack.

**Distributing the effort and knowhow of engineering the complex distributed system**, which can lead to a more transparent, well thought out, and standard way of realizing client-server MPC. Although MPC is distributed computing *per se*, it is often incompatible with foundational techniques of distributed systems e.g. computing parties having unproportional influence over the system due to the coordinator/master-slave pattern. Carbyne Stack, being an open source and community driven project, might help to reimagine and reimplement the measures we have so far invented for the application server with better qualities (reliability and resiliency).

Contributing to Carbyne Stack's aspiration for becoming an **industry standard for MPC** can be mutually beneficial: Sharemind MPC could adopt the unified API early on in the process to advance interoperability while refining Carbyne Stack's software requirements.

**Parallel task execution and scaling** comes naturally from how Carbyne Stack utilizes the nodes in a Kubernetes cluster to distribute MPC computation. It's observed, that task-level parallelization of MPC applications is a surefire way to improve execution times [3]. Unfortunately the current architecture of Sharemind MPC can not benefit from this in full potential, as concurrent tasks compete for network and compute resources. Both of these issues are helped by logical and physical distribution of tasks within a VCP's cluster.

**Augment Carbyne Stack with a range of Sharemind MPC features** like the Protocol DSL for protocol development and the SecreC language with its Analytics Library and standard library for writing optimized MPC programs. Further, this will enable previously developed Sharemind MPC applications to fully utilize modern environments such as commodity clouds.

**Seamless cloud deployment** procedures provide computing parties with tools and documentation to deploy a reproducible environment in commodity clouds. Carbyne Stack has a head start in this regard, with most cloud providers offering managed Kubernetes control planes, and cloud-specific IaC scripts are in the works. This has potential to reduce the assistance of set up needed from the technology provider to a minimum – something that hampers on-premises and bare metal deployments.

# 4 Integration Plan

---

## 4.1 Goals

The bottom line of the integration is to engineer a minimal composition of Sharemind MPC components into a fully functional MPC runtime that is operated in a similar fashion to the current MP-SPDZ binaries. The following are expectations for both the integration and the resulting architecture changes of Sharemind MPC.

1. The Sharemind MPC runtime should be interchangeable with the current MP-SPDZ variant. The backend choice should be a matter of a configuration change.
2. The new runtime comes in the form of an executable, spawnable by Ephemeral as it currently happens with MP-SPDZ.
3. The executable would contain the minimum to run a single MPC task (henceforth referred to as *task-as-process*) but retain the characteristic feature-set of Sharemind MPC, including SecreC, protocols and protocol DSL, previously written algorithms, etc.
4. The resulting architectural changes should not irrecoverably branch from the Sharemind MPC product. The task-as-process track is to eventually be a part of standalone Sharemind MPC along with any necessary accompanying services.

## 4.2 High Level Vision

The envisioned runtime is essentially a trimmed-down version of the Sharemind MPC application server. The trimming naturally begins from the functionality that Carbyne Stack already solves with its services (see section 2.1) or is expected to support in the future. The list of such responsibilities that are therefore transferred to Ephemeral and other services are as follows:

1. listening for clients' task invocations and executing the task;
2. task lifecycle management;
3. keeping tabs on and coordinating with neighbouring computing parties (VCPs);
4. handling of task arguments and published values;
5. client authentication\*;
6. client authorization, access and policy control\*\*;
7. securing protocol communication between task instances\*\*;
8. submission and storage of MPC programs\*\*;
9. guarantee transactionality of tabular secret storage facilities\*\*;
10. session management for a scheduled execution model\*\*.

Asterisk (\*) marks the feature which is currently worked on; double asterisks (\*\*) mark the features which do not yet have a proposal.

This would leave the runtime comprising two facilities:

1. Networking layer for protocol communication;
2. Virtual machine with the loadable module system.

## 4.3 Development Roadmap

Being in its early stages of development, many technical challenges still lie ahead for Carbyne Stack itself before being ready for production use. A few examples of this were already listed in section 4.2 (items 5, 6, 7, 8) – elements not exclusive to this integration but rather necessities for client-server MPC – are subject to being realized through enhancement proposals. We can't predict the outcome of these developments (although we plan on participating in these processes), so we don't focus on these aspects in this section.

The integration plan is presented as a roadmap of milestones, that are adjusted and added as Carbyne Stack and the integration itself matures. Further, this gives a better overview on what functionality we sacrifice and regain over the integration period.

Milestones with lesser impact or such that can be considered enhancements are listed in no particular order in appendix A.

### 4.3.1 Milestone 1: A Minimal Runtime

This milestone's development focus is on the Sharemind MPC side, with the goal of creating the Sharemind MPC task-as-process runtime. Its starting point is the application server; the end marks a PoC of a SecreC task that can be run by Ephemeral with only small changes.

The transition from the application server to a single process per task removes the need for much of the complexity which was specific to the service platform – isolation of concurrent sessions and stream multiplexing to list a few. It subsequently does not make effort to preserve most of the fundamental functionality of Sharemind MPC. This includes the distributed databases and task I/O, client controller along with all access control features, consensus streams, and multiple concurrent protection domains<sup>1</sup>. Hence during this milestone, supporting only such programs which take no arguments and publish no results is sufficient, even though it provides no realistic utility.

We assume the following capabilities from the runtime:

1. preprocess bytecode and initialize the virtual machine and protection domain(s);
2. establish connections to counterpart runtimes during initialization
  - a. without mutual authentication and transport layer security;
  - b. skipping client process negotiation and validation of bytecode checksum;
  - c. establish connections with all nodes of the clique;
  - d. an established connection is further only used by the protection domain;
3. independently run computation until finished.

Item 2d allows for major simplifications to the networking layer if MPC programs were to only support one protection domain at a time. Otherwise, multiplexing can not be eliminated.

The executable would require at least the following arguments:

1. IP and TCP port of each neighbouring party's task instance;
2. MPC program's bytecode.

Finally, a fork of the Ephemeral server should be created, implementing the generic MPCEngine interface that executes the new runtime. The set-up is illustrated on figures 1 and 2.

---

<sup>1</sup>Multiple protection domains refers to the feature of Sharemind MPC to be configured with several VC topologies, that can utilize different protocols within a single SecreC program.

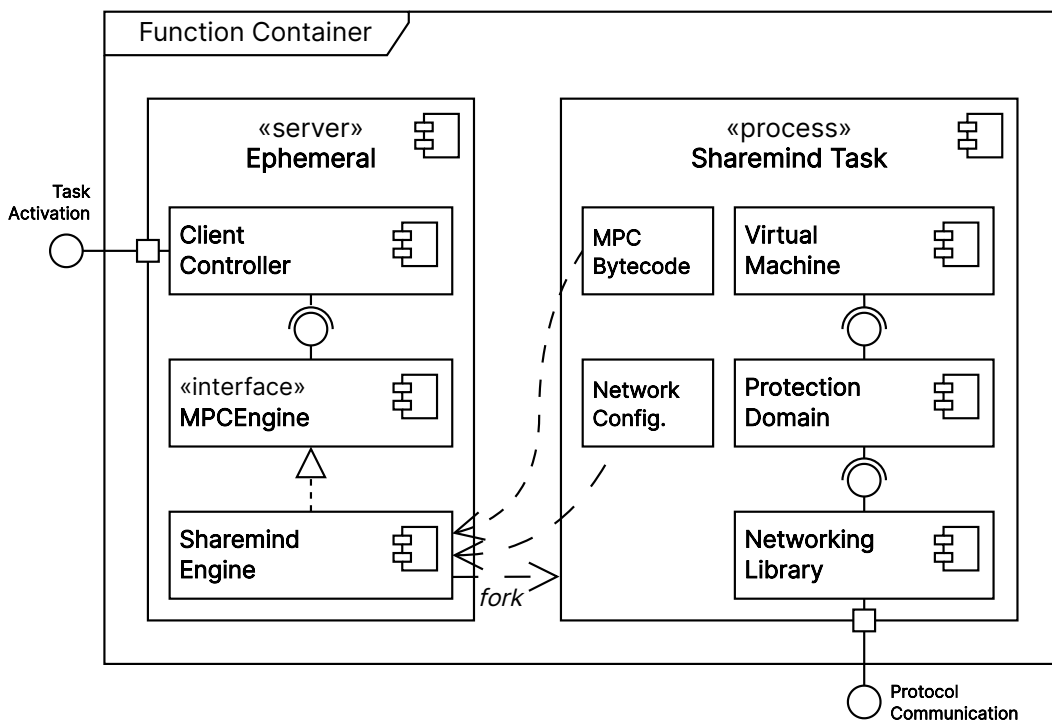


Figure 1. Component diagram showing the relation of the new runtime and Ephemeral

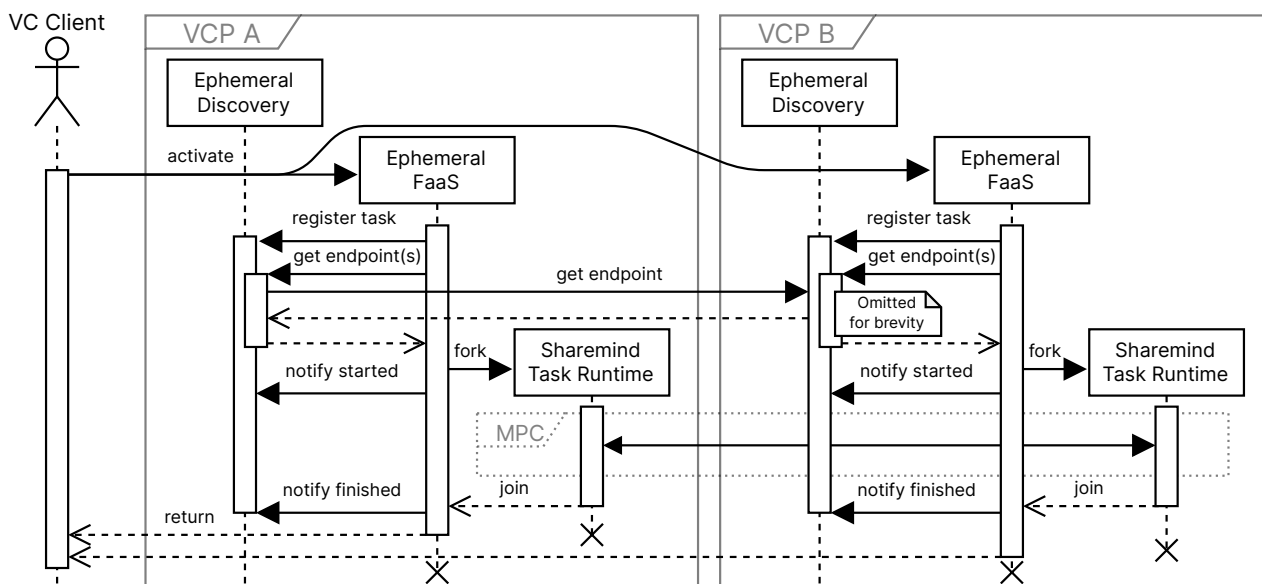


Figure 2. Sequence diagram illustrating the order of operations for invoking an MPC task within a two-party virtual cloud.

### 4.3.2 Milestone 2: Task Arguments and Published Values

This milestone targets both Carbyne Stack and the runtime to implement client supplied arguments and returned values.

SecreC programs rely on named arguments as inputs, typically transmitted from the client to the server as a binary argument stream along with the task execution command. The runtime has several options for acquiring these arguments outside of the application server, either reading them from

1. a file;
2. a Unix socket (currently used between Ephemeral and MP-SPDZ);
3. the standard input.

Each option is functionally equivalent. Published values, or the output, can follow the same path up to the calling program either by an intermediary file, socket, or standard output. As noted, the socket variant is used between Ephemeral and its current runtime, however it is different in a sense that the stream consists of a sequence of secrets, rather than named arguments. In Sharemind's case, the whole argument stream would need to be ingested fully, as the bytecode could use these values in any order.

The current state of the Ephemeral Client Library insists<sup>2</sup>, that inputs and outputs are envisioned to be mediated by Amphora. The extra step is likely necessary due to the fact that the Amphora client and server facilitate secret share authentication procedures required for active security. Since our primary target is our passive protocol that only works on plain additive secret shares, we propose to not support Amphora with the runtime for the time being. This is mainly due to the high development cost of modifying Amphora to support unauthenticated secret shares, public values, and the range of data types we need supported for SecreC.

The work that needs to be carried out on Carbyne Stack's side is as follows:

1. extend the Ephemeral Client Library to
  - a. accept a binary arguments object and pass it along with the activation request;
  - b. handle the response which may contain a binary published values object.
2. supplement the `MPCEngine` implementation created in milestone 1 to
  - a. provide the binary argument object to the runtime using the decided medium (file, socket, or standard input);
  - b. write the resulting binary object of published values to the HTTP response.

### 4.3.3 Milestone 3: Distributed Tabular Databases

Sharemind MPC owes much of its performance to vectorized operations on secret data to batch together protocol messages. Specific language constructs in SecreC enable writing code in a SIMD manner, which is highly encouraged over processing scalars and using loops. This has also influenced the persistent storage method of choice in Sharemind MPC to be columnar databases with its `tabledb` module on which numerous applications, Rmind among others, are dependent on. With the interest of ensuring backwards compatibility, accessing tabular data within MPC programs is required.

---

<sup>2</sup><https://github.com/carbynestack/ephemeral/blob/a39a1b2e8cb2e8a9be546b17544573db15c86630/ephemeral-java-client/src/main/java/io/carbynestack/ephemeral/client/Activation.java#L24>

Application server's current database engine relies on files in the HDF5 file format to store the data. An issue with this approach in a Kubernetes cluster is the absence of a common filesystem without Persistent Volumes with a ReadWriteMany access mode. If using a cloud file storage solution like Amazon's EFS, this solution becomes non-portable, while in-cluster file storage solutions like Rook Ceph are heavyweight and store data in temporary node storage [3].

In search of a cloud native solution, some possible approaches are as follows.

1. **An Amphora approach** would use the key-value storage to create an abstraction of a columnar database. First, it requires the same prerequisite changes as mentioned in milestone 2 for the Amphora database. Further, the database values would have to support array types to emulate columns. A `tabledb` module would have to be created for interfacing with Amphora within the runtime. To complete the table abstraction, there would have to be storage interface for metadata, which collates the column vectors. A foreseeable sacrifice would be not having a possibility to modify or append to tables due to the nature of object storage.
2. **Dedicated database approach** involves adopting an existing cloud native DBMS to back secret storage. Settling on the exact technology of choice depends on a variety of factors and requires in-depth analysis.

It's important, that growable concurrent-access data tables need a custom layer of coordination on top of any off-the-shelf DBMS to guarantee consistency over a VC. The current application server applies a *consensus* protocol to ensure correct ordering of writes [4]. This or an analogous component will have to be implemented as a cluster-wide service.

The tabular database support would need to be followed up with ACL enforcement, the specifics of which depend on Carbyne Stack's future authorization facilities.

# Bibliography

---

- [1] The Carbyne Stack Authors. *Carbyne Stack*. <https://carbynestack.io>, Last accessed on 2022-12-30. 2022.
- [2] Marcel Keller. "MP-SPDZ: A Versatile Framework for Multi-Party Computation". In: *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*. 2020. DOI: 10.1145/3372297.3417872. URL: <https://doi.org/10.1145/3372297.3417872>.
- [3] Kert Tali. "Parallel and Cloud-Native Secure Multi-Party Computation". [https://comserv.cs.ut.ee/ati\\_thesis/datasheet.php?id=74954](https://comserv.cs.ut.ee/ati_thesis/datasheet.php?id=74954). MA thesis. Institute of Computer Science, University of Tartu, 2022.
- [4] Cybernetica AS. *Deliverable D3.2: A protocol suite for robust distributed computations*. SEVILLA Technical document. 2016.

# Appendix A Extra Milestones

---

## A.1 Task Scheduling

Client-server MPC is often prone to failures due to intermittent network faults between the client and the server if the client has no way to resume a session. This issue is especially apparent with long-running MPC tasks that process large data volumes. We find that a task scheduling API, the likes of Hadoop Yarn or other big data processing platforms, is more fitting to client-server MPC than the current synchronous model used by both Carbyne Stack and Sharemind MPC.

## A.2 Transactional Distributed Databases

As described in section 2.1 and milestone 3, tasks in Sharemind MPC have write access to databases during execution. Databases are pseudo-replicated in a manner in which every VCPs database is a mirror image of others regarding schema and shape, with each private data value representing its secret share. This evidently presents the requirement for replicas to maintain a strongly consistent state – all parties *not* reading a share corresponding to the same value or receiving an error would constitute a wrong result or failure.

The problem arises with concurrently accessed structured databases. Ephemeral and Amphora effectively avoid this issue by using client specified unique keys for each new value. With Sharemind MPC only the correct ordering of database operations is guaranteed. ACID compliant transactions and a strong consistency guarantee in secret share databases are necessary going forward. This is to avoid manual intervention if databases were to end up in an inconsistent state, which is especially laborous in an MPC set-up.

We suggest transactions to be atomic over a task – this makes sense in that there is no way for tasks to be checkpointed or resumed. This would enable full transparency for the developer and compability with legacy SecreC code.

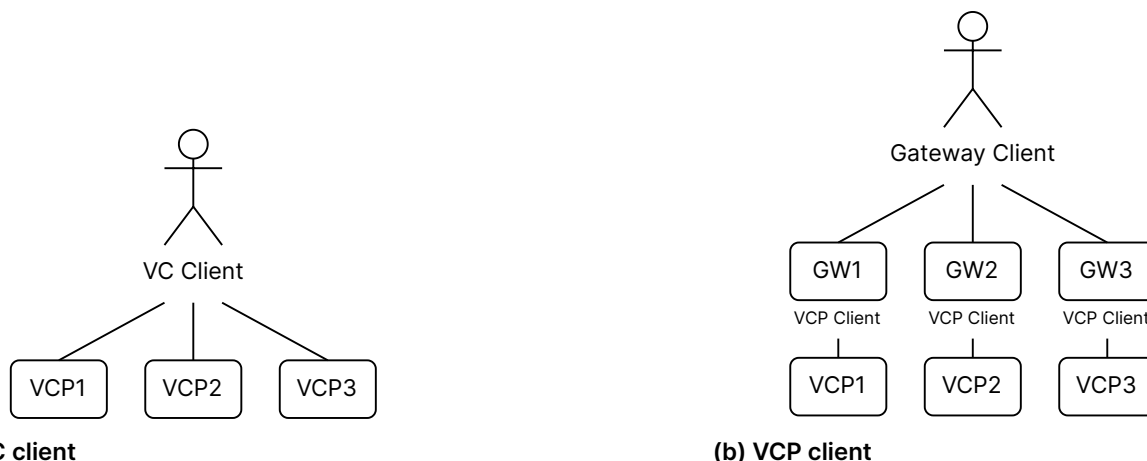
While the details of the transactional architecture are still in a draft stage, the specific isolation level is an open problem. It is however clear, that all commits to the distributed database have take place under a global lock to ensure no discrepant reads in concurrently running tasks. The consensus service mentioned in milestone 3 could be supplemented with lock management and any further ACID measures.

## A.3 Gateway Applications

A distinction can be made between cases where the clients are data owners or data receivers who directly trigger MPC tasks, or an application initiating on its own behalf, possibly containing some public business logic. In Sharemind MPC, these applications are called (web) gateways, built on a supplied gateway library. The benefit of these gateways has so far mostly been the opportunity to integrate MPC subroutines into web applications. In this case the gateway translates HTTP from the web client to the binary over raw TCP protocol used by the application server, that can't be acommodated by a web browser. While Carbyne Stack has a native HTTP client API, support for extra server-side business logic can still be beneficial.

For clarity let's call the classic model of client, that communicates and synchronizes operations with all parties, a VC client (figure 3a), and a client that only communicates to a specific party





(a) VC client

(b) VCP client

**Figure 3. Client models for interacting with the MPC platform.**

a VCP client (figure 3b). Each gateway (GW) acting as the VCP client is usually in the same administrative domain with its VCP.

This type of deployment offers flexibility in realizing various business cases. For example, each server could buffer input and/or result shares, schedule the execution to a certain time (epoch), pipeline several tasks, coordinate parallel tasks. The VCP client method enables MPC powered applications which ingest secret shares over any medium, like streams or web clients, utilizing any authentication, without the need for a central service and therefore avoid encrypting shares for VCPs.<sup>1</sup>

A motivating example is the Sharemind Secure Survey, in which subjects receive a token to answer a set of questions on a distributed web application. Upon submission, the answers are secret shared in the browser among three gateways, which immediately triggers an MPC task to validate that multiple choice questions are within the allowed range. The gateways therefore play a role in enforcing application specific rules in effort to avoid malicious clients poisoning the dataset. As the survey closes, the three gateways trigger the second MPC task to perform analysis and store the aggregate results for the survey owner to pull and recombine on demand.

<sup>1</sup>An alternative approach is to use a central, or single gateway with a VC client to mediate shares encrypted for each VCP with a respective public key. This requires an intricate key distribution to be in place. Further the central service has to be trusted to not selectively omit inputs from the computation.