

**RADU POPESCU**

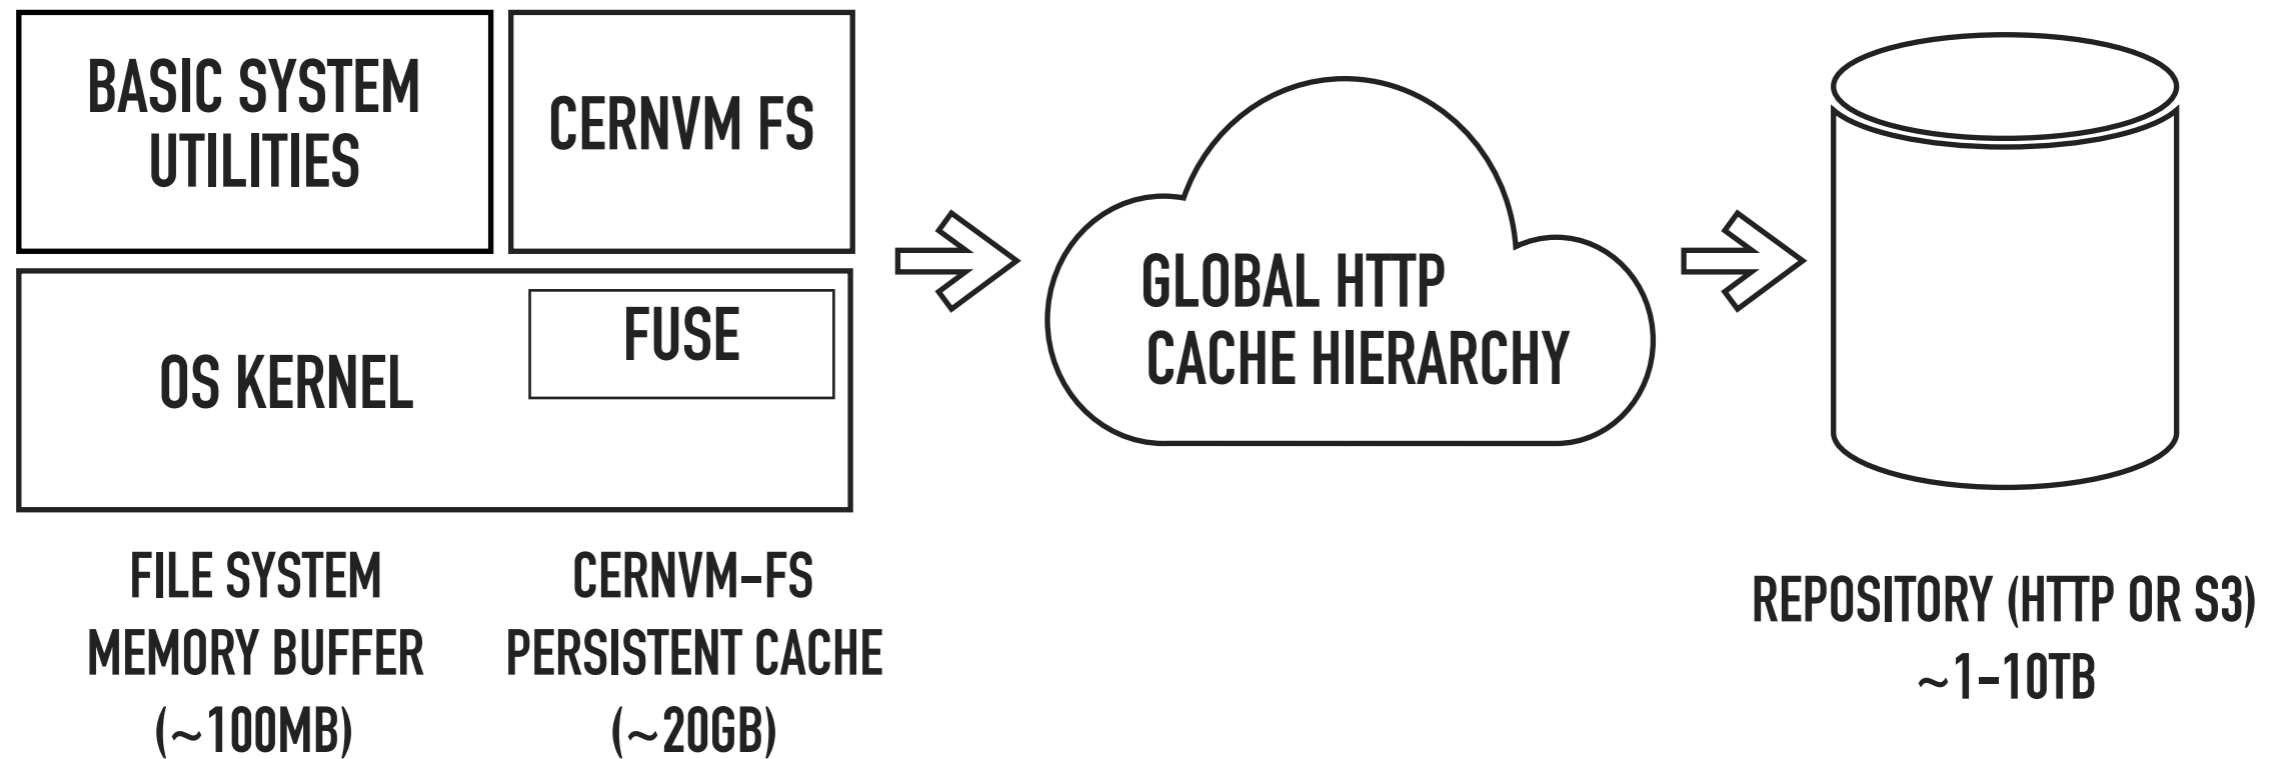
CERN EP-SFT, ON BEHALF OF THE CERNVM-FS TEAM

---

# **ADDING HTTP INGESTION SUPPORT TO XCACHE**

**XROOTD@IN2P3 LYON, JUNE 2019**

# CVMFS – A FILE SYSTEM APPROACH TO DISTRIBUTING SOFTWARE



- FUSE based, independent mount points, e. g. /cvmfs/sft.cern.ch
- Clients have a read-only view; single writer into repository
- Immutable, content-addressed storage
- HTTP transport, access and caching on demand

---

## CVMFS HTTP CACHING LAYER

- Essential to CVMFS scalability at global scale
- Based on Squid (<http://www.squid-cache.org/>)
  - Single-threaded, performance for large files?
- Separate from existing caching infrastructures / CDNs

## PREMISE

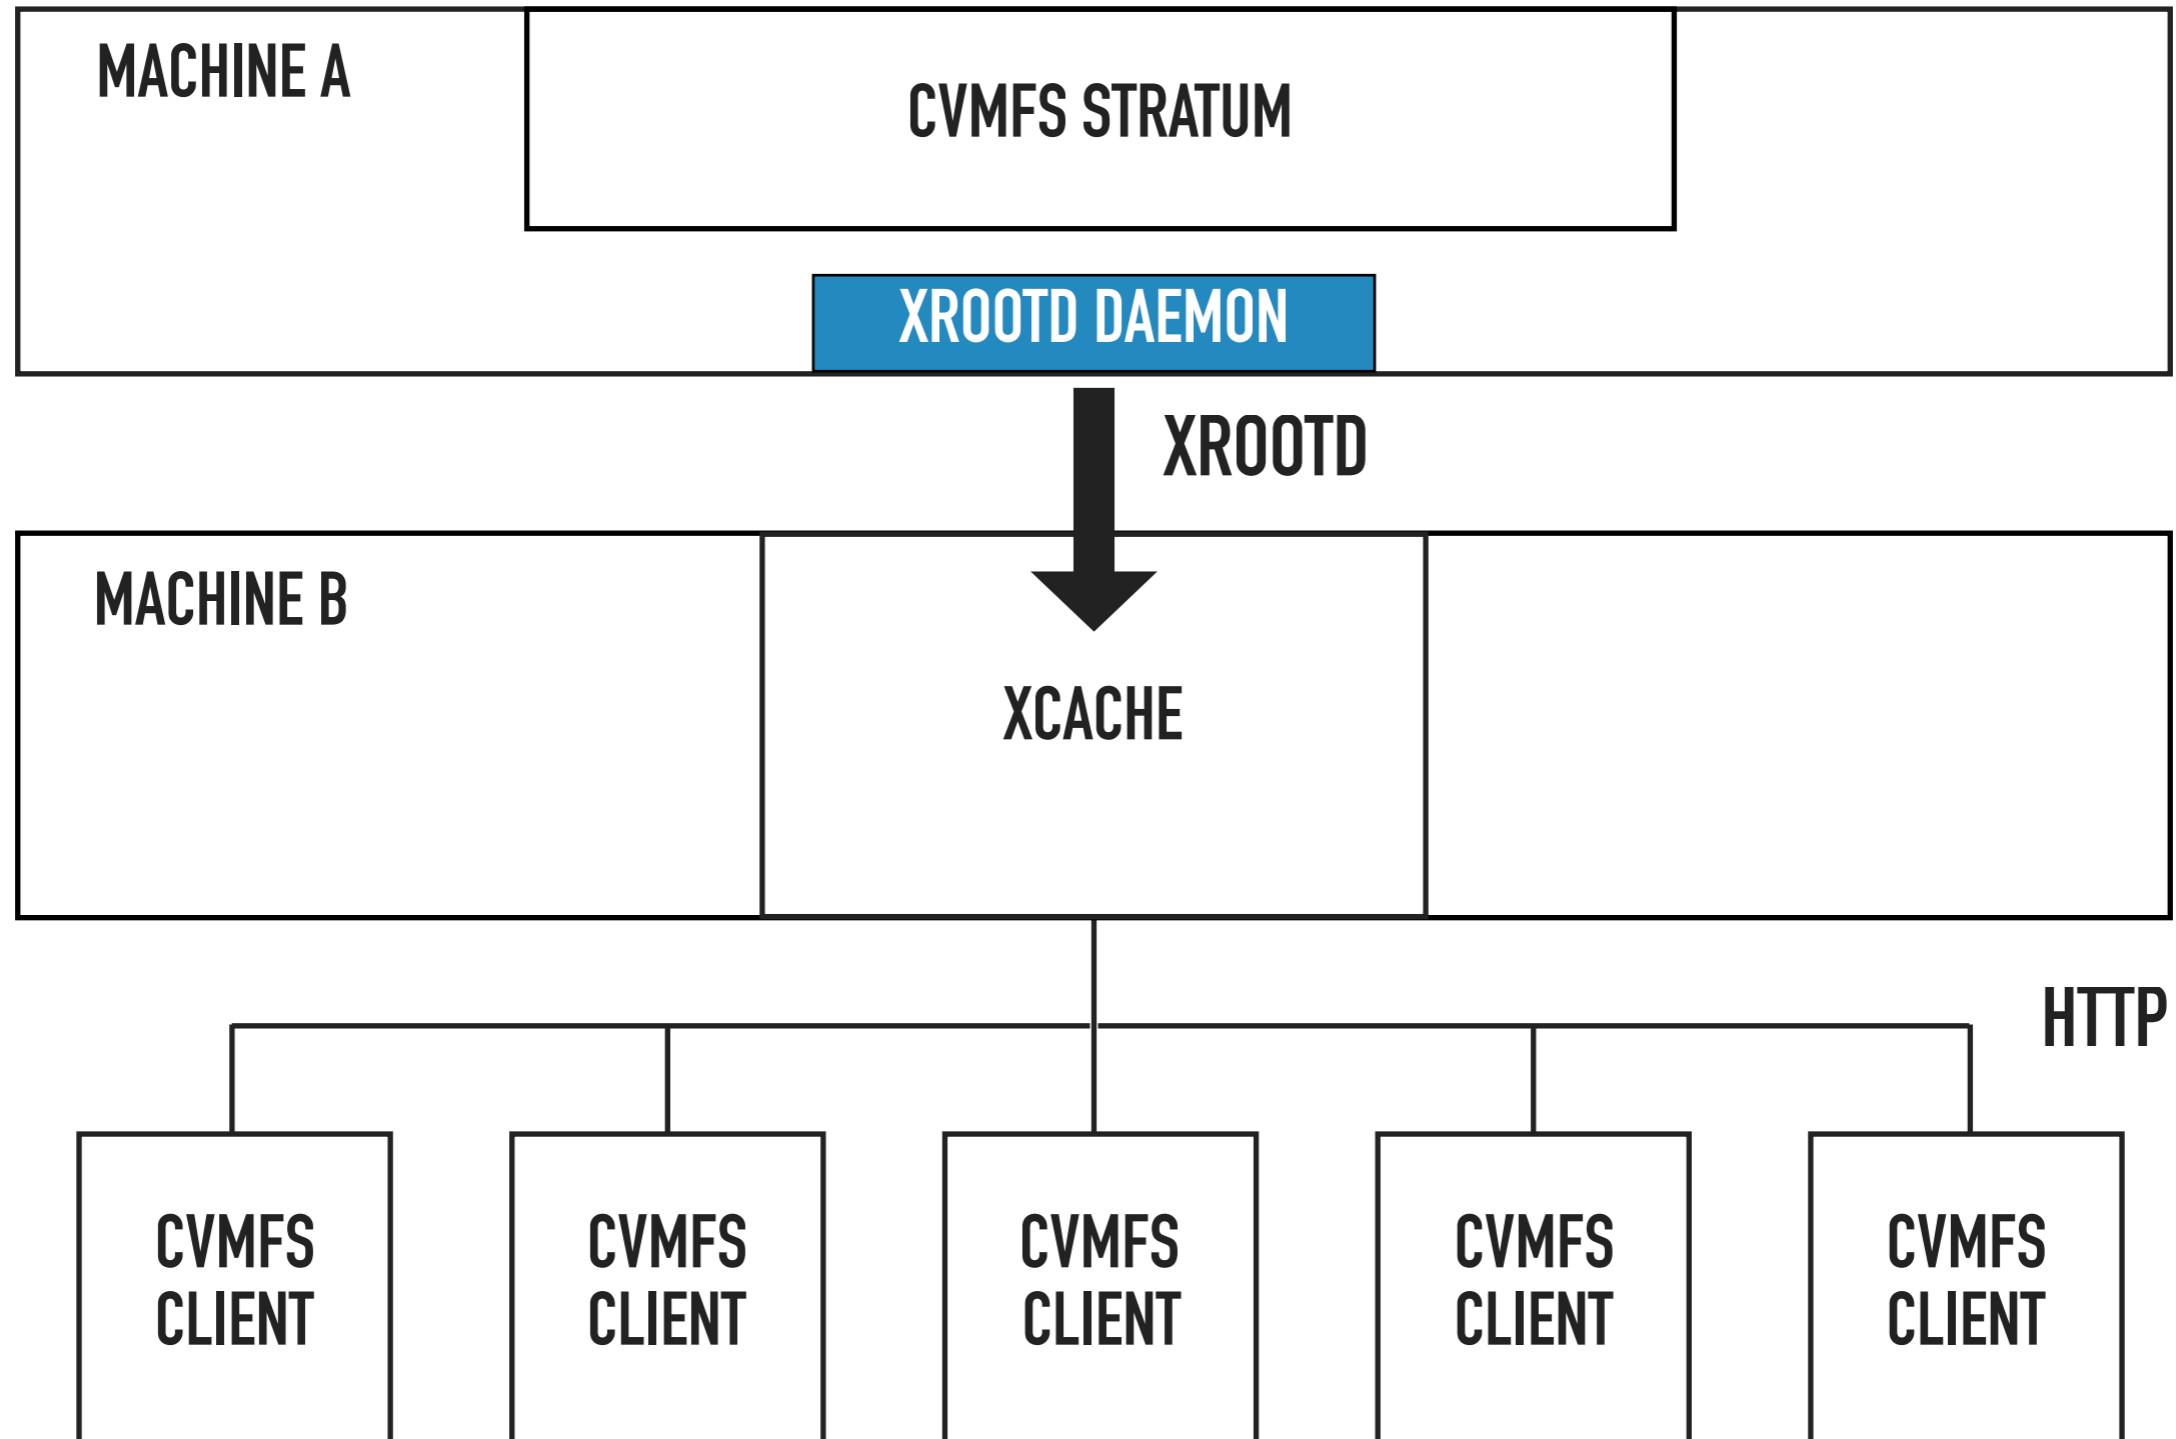
- Can XRootD be used as an alternative transport for CVMFS?
- An XCache instance can be introduced without changes to the CVMFS clients' source

Not the first time we look at XRootD: XRootD-based cache plugin for CVMFS (2017): <https://github.com/cvmfs/xrootd-plugin>

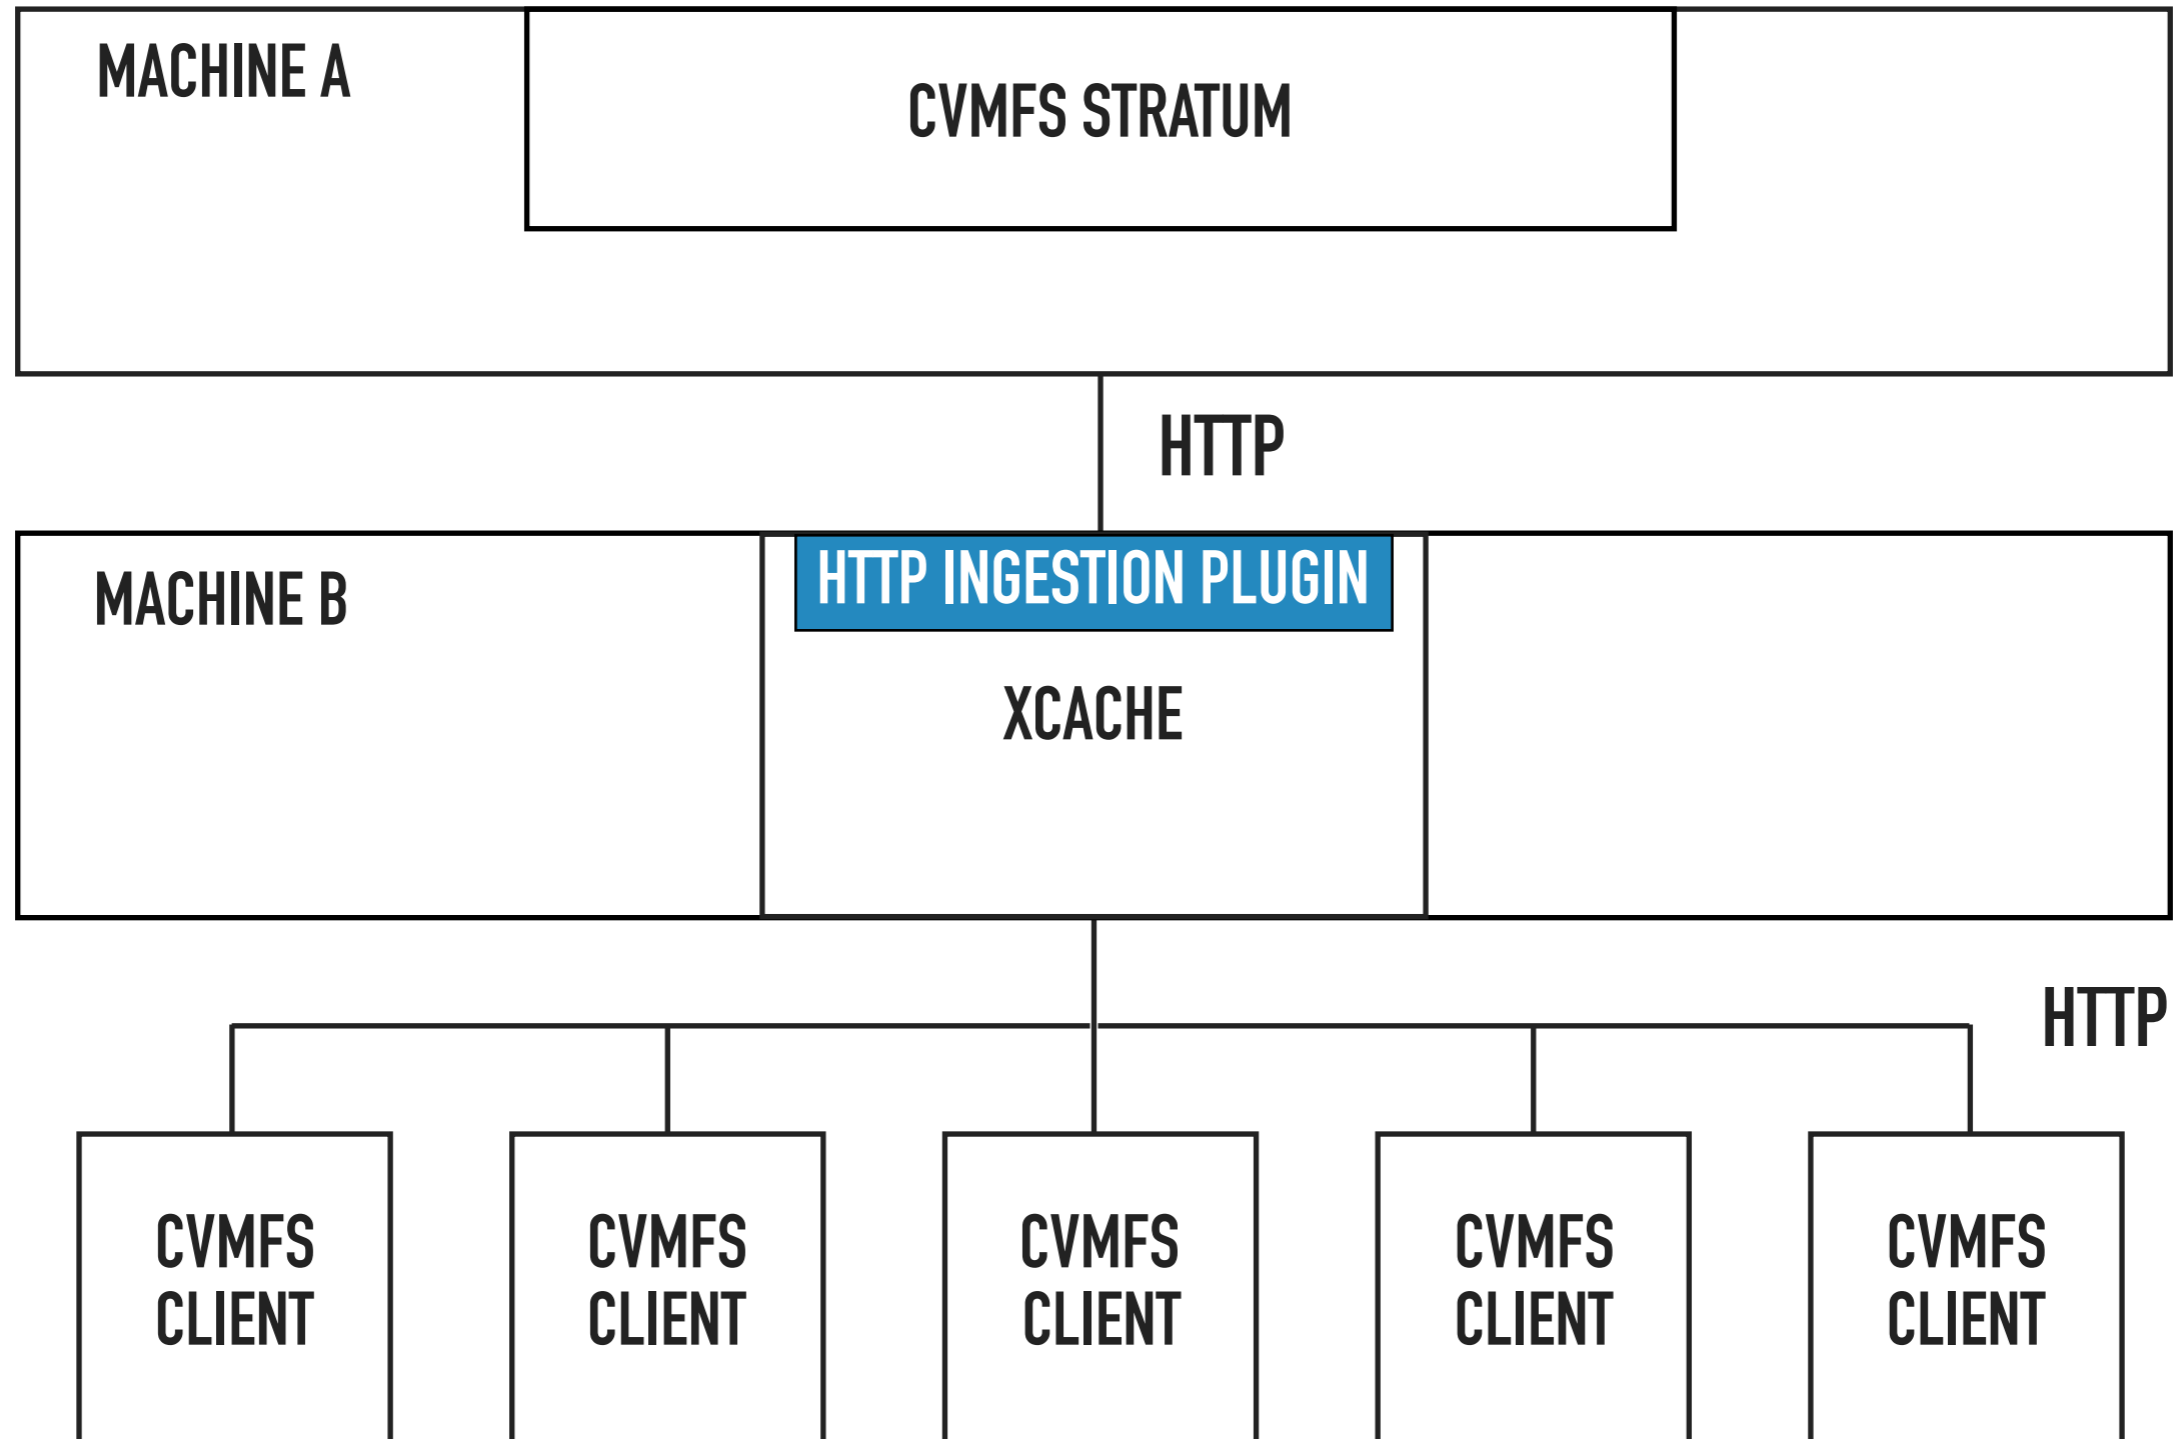
## XROOTD/XCACHE ADVANTAGES

- Scalable via horizontal clustering
- Multi-threaded
- Flexible
- Self managing in terms of disk utilization
- Architected for local or regional use
- Accessible via HTTP[S] or XRoot protocols

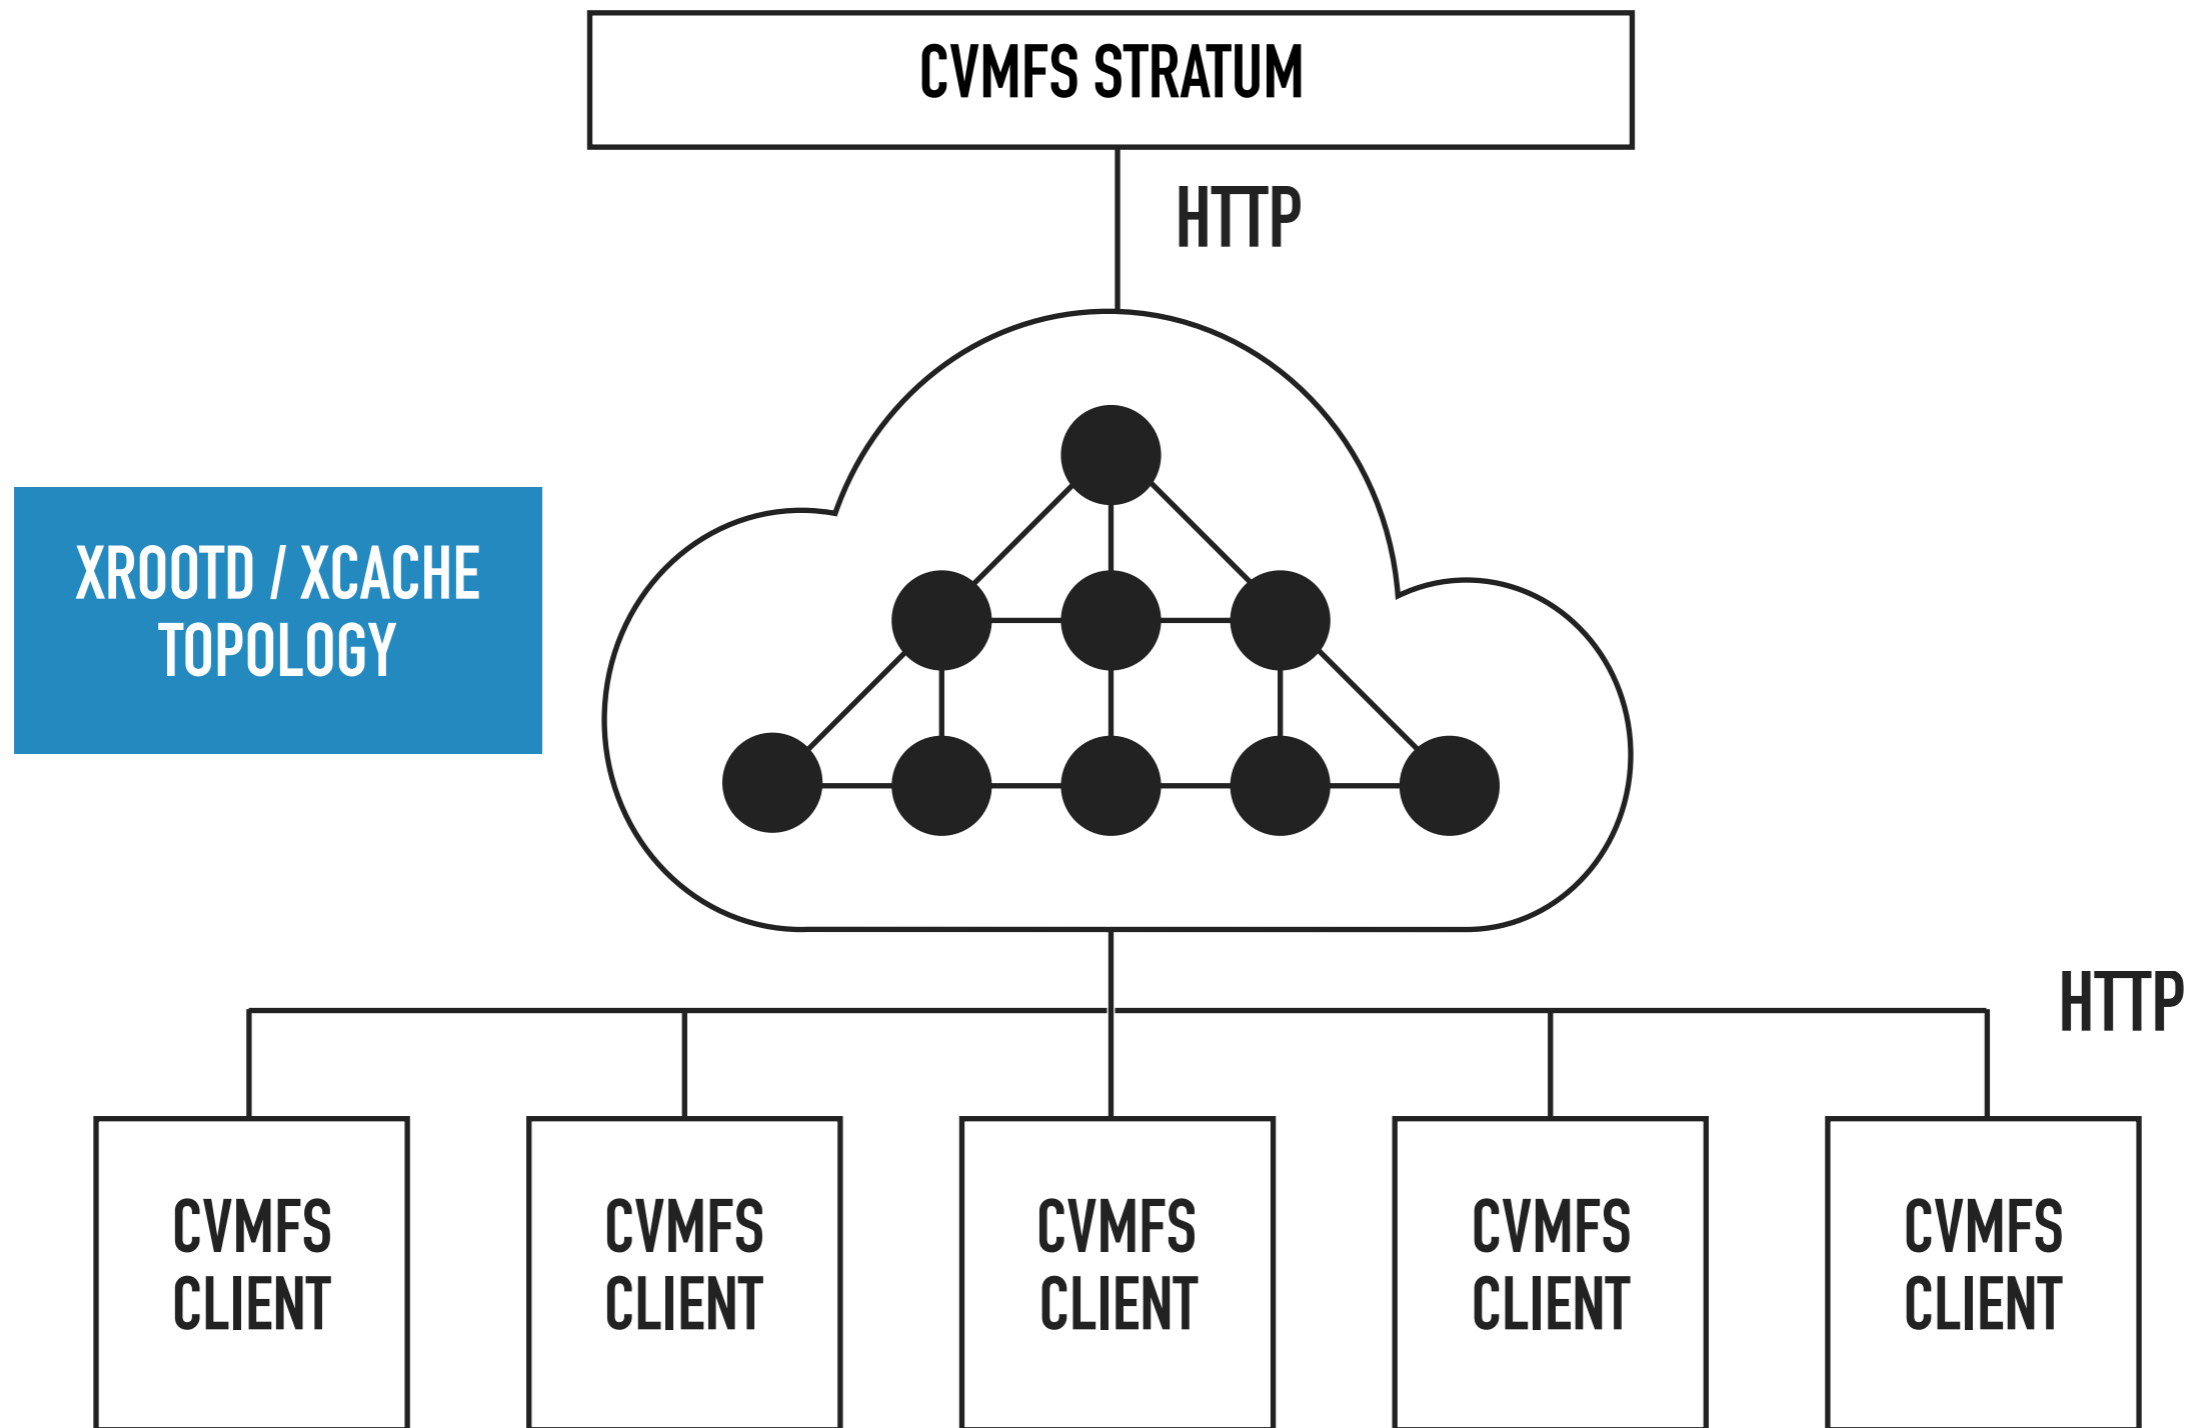
# INITIAL TEST DEPLOYMENT



# IMPROVED DEPLOYMENT



# POTENTIAL SETUP





## INTRODUCING: XRDCL-HTTP

- Allows the XRootD client to read from HTTP sources
- Some support for writing to HTTP locations
- Based on XrdClPluginInterface (FilePlugin, FileSystemPlugin)
- Uses Davix
- Prototype, not production-quality

---

## XROOTD CLIENT PLUGIN INTERFACE

- Almost all operations with meaning in the context of HTTP are implemented
- Not implemented: Sync (NoOp), Fcntl, Visa, Locate, Truncate, Chmod, Ping, DirList, StatVFS, Protocol, SendInfo, Prepare
- Based on XrdCl plugin interface: can be used with XCache and xrdcp!

## XCACHE CONFIGURATION

- `libXrdClHttp-4.so` should be available in client plugin search path (set `XRD_PLUGINROOT` if needed)

- XrdCl plugin configuration:

```
url = http://*
```

```
lib = <<XROOTD_PREFIX>>/lib/libXrdClHttp.so
```

```
enable = true
```

- Set the XCache origin to the HTTP URL:

```
pss.origin http://cvmfs-stratum-zero-hpc.cern.ch:80
```

## XRDCP CONFIGURATION

- libXrdClHttp-4.so should be available in client plugin search path (set XRD\_PLUGINROOT if needed)
- Client plugin configuration
- New flag (-A) enables HTTP sources and destinations:

```
xrdcp -A \
```

```
http://some.host.com:8080/file.txt \
```

```
xroot://localhost:1095///destdir/file.txt
```

## XRDCP CONFIGURATION

- libXrdClHttp-4.so should be available in client plugin search path (set XRD\_PLUGINROOT if needed)
- Client plugin configuration
- New flag (-A) enables HTTP sources and destinations:

```
xrdcp -A \
```

```
  xroot://localhost:1095///srcdir/file.txt \
```

```
  http://some.host.com:8080/file.txt
```

## XRDCP CONFIGURATION

- libXrdClHttp-4.so should be available in client plugin search path (set XRD\_PLUGINROOT if needed)
- Client plugin configuration
- New flag (-A) enables HTTP sources and destinations:

```
xrdcp -A \  
  
```

```
  http://some.host.com/file.txt \  
  
```

```
  http://another.host.com///file.txt
```

# DAVIX

- Library for file management over HTTP protocol
- Using the POSIX interface of Davix
- Single-threaded (delegate parallelism to XRootD)
- Straightforward use for most operations
- Varying levels of support for certain operations (nested directories?) depending on HTTP or WebDAV

---

## TESTING XRDCL-HTTP

- Basic functional tests:
  - Create random files and dirs
  - Caddy serves the files over HTTP
  - xrdcp with xrdcl-http accesses and modifies the HTTP source
- It would be nice to have a more structured approach:  
QuickCheck, Hypothesis?
- Check WebDAV, HTTP, etc...



## BENCHMARK SETUP

- Randomly generate 100MB worth of files: 10 x 10MB, 100 x 1MB, 1000 x 100kB
- Served locally with Caddy HTTP server and XRootD
- Download all files with xrdcp, in turn, from HTTP and XRootD sources

# BENCHMARK RESULTS

## Average throughput (MB/s)

Configuration	HTTP	XRootD
10MB x 10	178.8	183.1
1MB x 100	33.2	33.1
100kB x 1000	3.6	3.5

# SUMMARY

- A new plugin for XCache and xrdcp allows ingesting from HTTP sources
- Distributed with XRootD 4.9
- Needs advanced test suite (QuickCheck, Hypothesis?)
- Use cases beyond CVMFS?
- Available at: <https://github.com/xrootd/xrdcl-http>

**THANK YOU!**