



# Multitier Modules in ScalaLoci

---

Pascal Weisenburger, Guido Salvaneschi  
TU Darmstadt, Germany

# Programming Distributed Systems

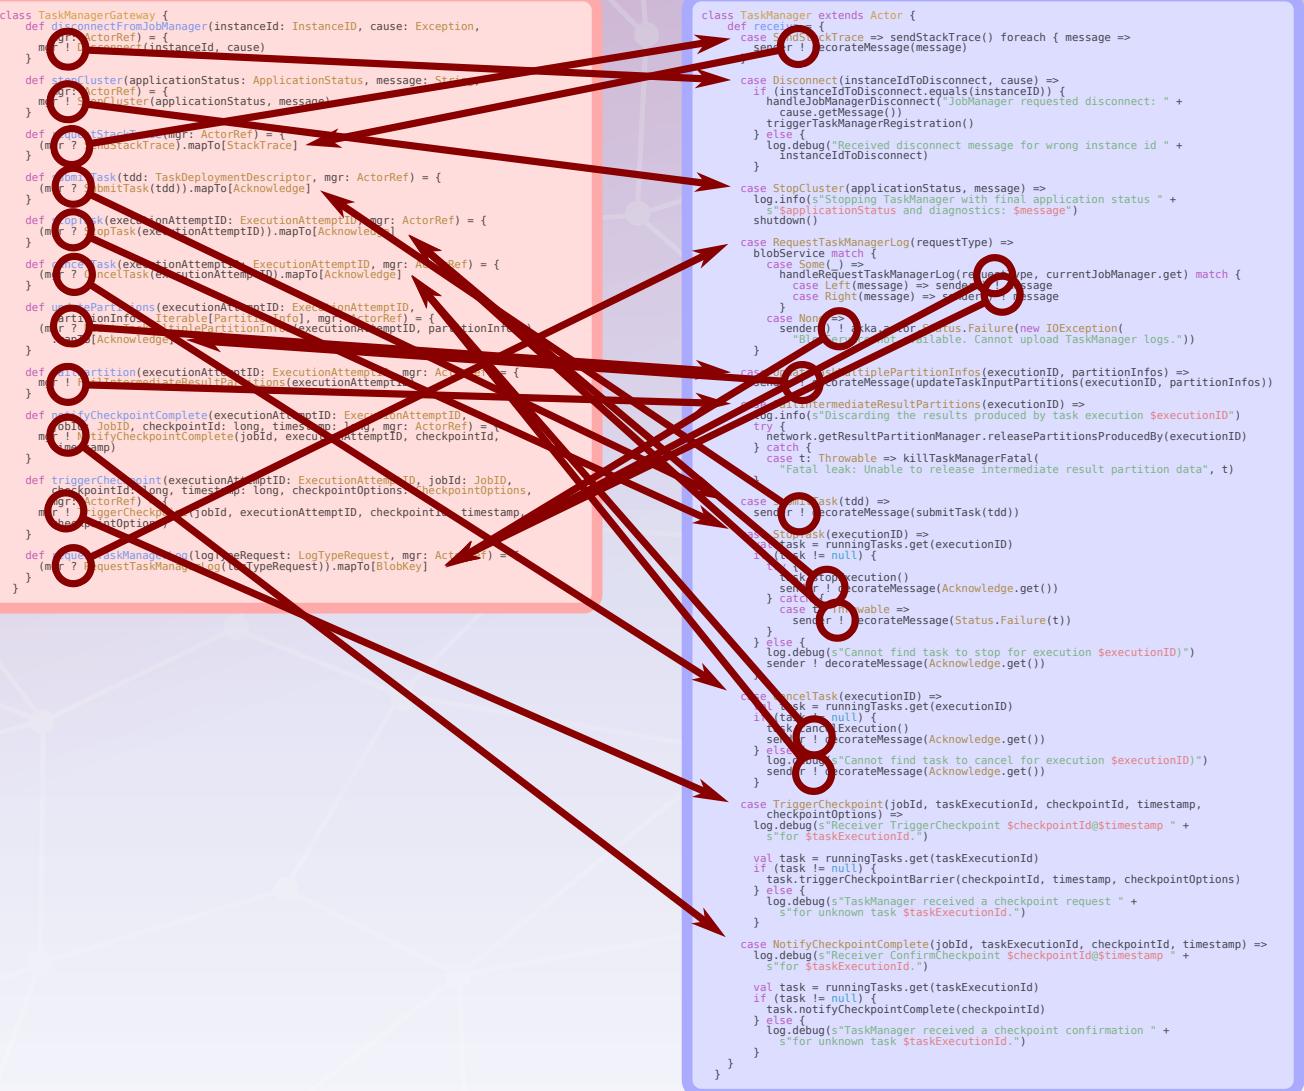
---

Developing distributed systems is *hard*

- Consistency
- Replication
- Fault Tolerance
- Distributed functionalities and communication

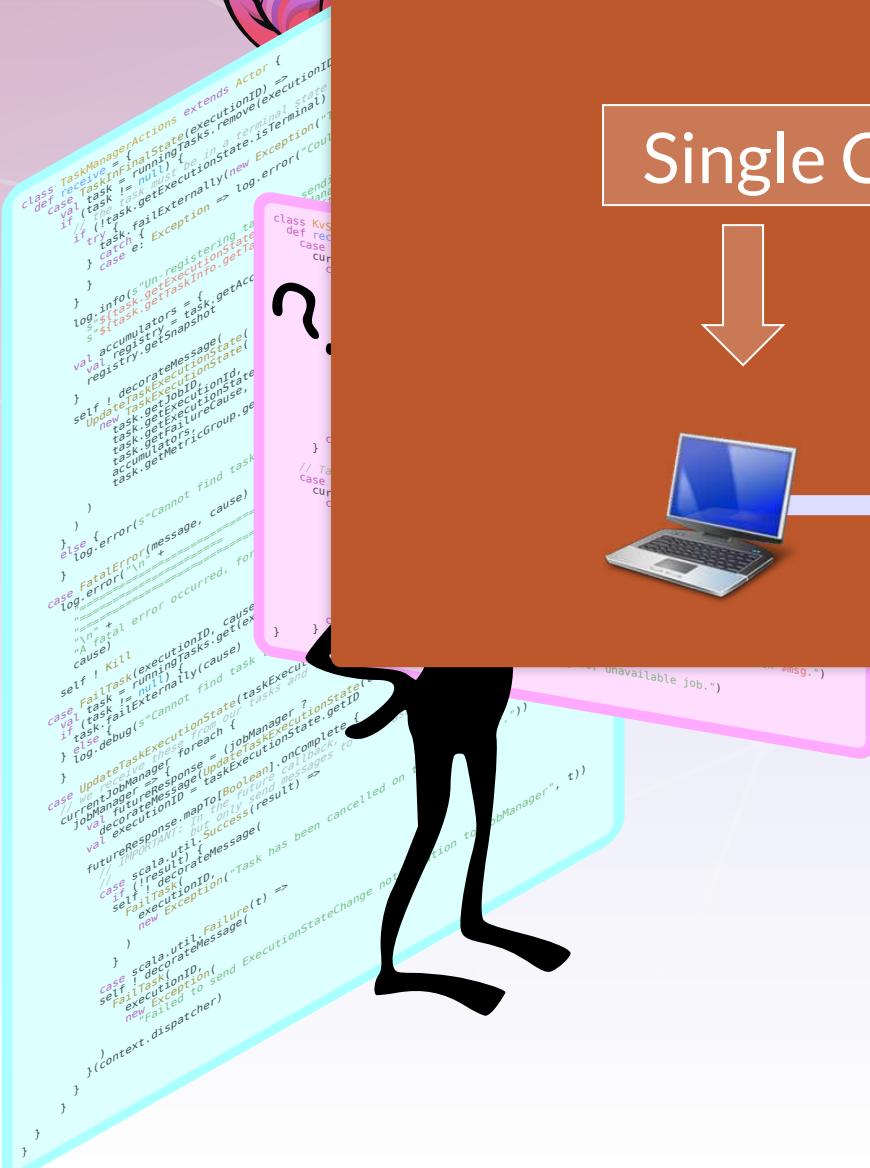


# Flink



# Multitier Languages

## Single Compilation Unit

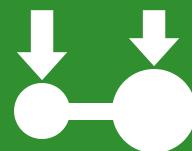


# ScalaLoci

---



Generic Distributed Architectures



Placement Types

# Placement Types

```
@peer type Master  
@peer type Worker
```

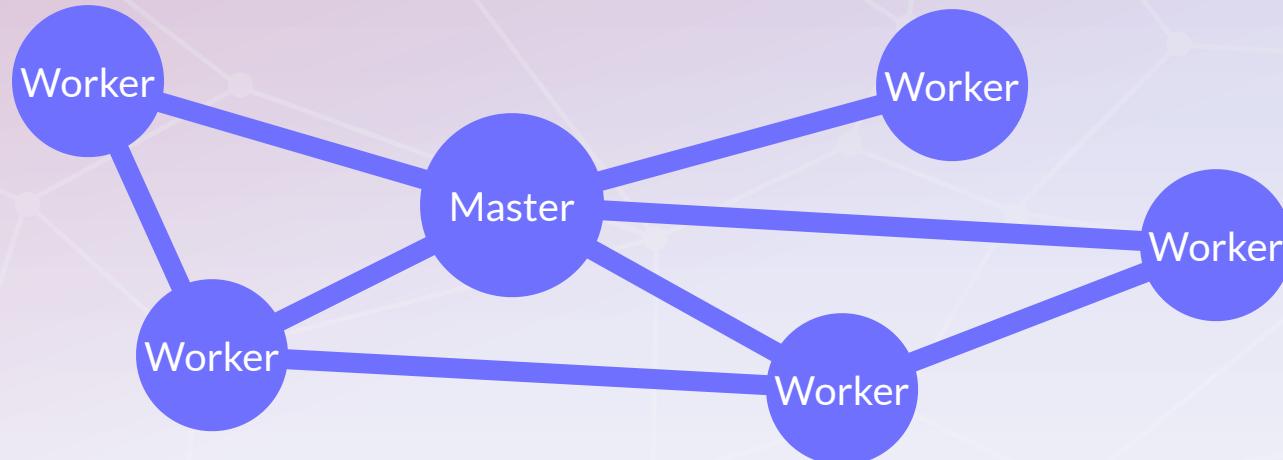
Peers

```
val tasks: List[Task] on Master  
= placed { getTaskList() }
```

Placement  
Types

# Architecture

```
@peer type Master { type Tie <: Multiple[Worker] }
@peer type Worker { type Tie <: Single[Master] with Multiple[Worker] }
```



Architecture Specification  
through Peer Types



Eliminated 23 non-exhaustive pattern matches  
and 8 type casts

```
class TaskManagerGateway {
    def receiveFromJobManager(instanceID: InstanceID, cause: Exception,
        msg: InstanceID) = {
        log.info(s"Received disconnect message for wrong instance id $instanceID")
    }
    def handleCluster(applicationStatus: ApplicationStatus, message: String) = {
        log.info(s"Received cluster status update: $applicationStatus")
    }
    def handleTask(tdd: TaskDeploymentDescriptor, mgr: ActorRef) = {
        (m: ?) match {
            case TaskDeploymentDescriptor(mgrRef) => mapTo[mgrRef]
            case _ => mapTo[StackTrace]
        }
    }
    def handleTaskAttempt(tdd: TaskDeploymentDescriptor, attemptID: ExecutionAttemptID, mgr: ActorRef) = {
        (m: ?) match {
            case TaskDeploymentDescriptor(mgrRef) => mapTo[mgrRef]
            case _ => mapTo[StackTrace]
        }
    }
    def handlePartitions(executionAttemptID: ExecutionAttemptID, mgr: ActorRef) = {
        (m: ?) match {
            case TaskDeploymentDescriptor(mgrRef) => mapTo[mgrRef]
            case _ => mapTo[StackTrace]
        }
    }
    def handleTask(executionAttemptID: ExecutionAttemptID, mgr: ActorRef) = {
        (m: ?) match {
            case TaskDeploymentDescriptor(mgrRef) => mapTo[mgrRef]
            case _ => mapTo[StackTrace]
        }
    }
    def handleTaskFailure(executionAttemptID: ExecutionAttemptID, mgr: ActorRef) = {
        (m: ?) match {
            case TaskDeploymentDescriptor(mgrRef) => mapTo[mgrRef]
            case _ => mapTo[StackTrace]
        }
    }
}
```

```
class TaskManager extends Actor {
    def receive = {
        case StopCluster(applicationStatus: ApplicationStatus, message: String) =>
            log.info(s"Stopping TaskManager with final application status $applicationStatus and diagnostics: $message")
            shutdown()
        case RequestTaskManagerLog(requestType) =>
            blobService match {
                case Left(message) => sender ! message
                case Right(message) => sender ! message
            }
            case None => sender ! Status.Failure(new IOException("No blob storage available. Cannot upload TaskManager logs."))
        case SubmitMultiplePartitionInfos(executionID, partitionInfos) =>
            partitionInfos.foreach { info => updateTaskInputPartitions(executionID, partitionInfos) }
        case NotifyCheckpointComplete(executionAttemptID: ExecutionAttemptID, jobID: JobID, checkpointID: Long, times: Long, mgr: ActorRef) =>
            log.info(s"Discarding the results produced by task execution $executionAttemptID")
            try {
                network.getResultPartitionManager.releasePartitionsProducedBy(executionID)
            } catch {
                case t: Throwable => killTaskManagerFatal(
                    s"Fatal leak: Unable to release intermediate result partition data", t)
            }
        case TriggerCheckpoint(executionAttemptID: ExecutionAttemptID, jobID: JobID, checkpointOptions: CheckpointOptions, timestamp: Long) =>
            val task = runningTasks.get(taskExecutionId)
            if (task != null) {
                task.triggerCheckpointBarrier(checkpointID, timestamp, checkpointOptions)
            } else {
                log.debug(s"TaskManager received a checkpoint request " +
                    s"for unknown task $taskExecutionId")
            }
        case NotifyCheckpointComplete(jobID, taskExecutionId, checkpointID, timestamp) =>
            log.debug(s"Receiver ConfirmedCheckpoint $checkpointID@$timestamp " +
                s"for $taskExecutionId")
            val task = runningTasks.get(taskExecutionId)
            if (task != null) {
                task.notifyCheckpointComplete(checkpointID)
            } else {
                log.debug(s"TaskManager received a checkpoint confirmation " +
                    s"for unknown task $taskExecutionId")
            }
        case Disconnect(instanceIDToDisconnect, cause) =>
            if (instanceIDToDisconnect.equals(instanceID)) {
                handleJobManagerDisconnect("JobManager requested disconnect: " + cause)
            } else {
                log.debug(s"Received disconnect message for wrong instance id " + instanceIDToDisconnect)
            }
        case StopCluster(applicationStatus: ApplicationStatus, message: String) =>
            log.info(s"Stopping TaskManager with final application status $applicationStatus and diagnostics: $message")
            shutdown()
        case RequestStackTrace(mgr: Remote[TaskManager]) =>
            on[JobManager] {
                case Left(stackTrace) => mgr ! stackTrace
                case Right(stackTrace) => mgr ! stackTrace
            }
            local.map(_.left.get)
        case SubmitTask(tdd: TaskDeploymentDescriptor) =>
            on[JobManager] {
                case Left(runnableTask) => runnableTask(tdd)
                case Right(stackTrace) => stackTrace
            }
            local.map(_.left.get)
        case StopTask(executionAttemptID: ExecutionAttemptID, mgr: Remote[TaskManager]) =>
            val task = runningTasks.get(executionAttemptID)
            if (task != null) {
                task.stopExecution()
                Left(Acknowledge.get())
            } else {
                log.debug(s"Cannot find task to stop for execution $executionAttemptID")
            }
            local.map(_.left.get)
        case CancelTask(executionAttemptID: ExecutionAttemptID, mgr: Remote[TaskManager]) =>
            val task = runningTasks.get(executionAttemptID)
            if (task != null) {
                task.cancelExecution()
                Left(Acknowledge.get())
            } else {
                log.debug(s"Cannot find task to cancel for execution $executionAttemptID")
            }
            local.map(_.left.get)
        case UpdatePartitions(executionAttemptID: ExecutionAttemptID, partitionInfos: java.util.List[PartitionInfo], mgr: Remote[TaskManager]) =>
            partitionInfos.foreach { info => updateTaskInputPartitions(executionAttemptID, partitionInfos) }
            local.map(_.left.get)
        case FailPartition(executionAttemptID: ExecutionAttemptID, mgr: Remote[TaskManager]) =>
            log.info(s"Discarding the results produced by task execution $executionAttemptID")
            network.getResultPartitionManager.releasePartitionsProducedBy(executionAttemptID)
            catch {
                case t: Throwable => killTaskManagerFatal(
                    s"Fatal leak: Unable to release intermediate result partition data", t)
            }
        case NotifyCheckpointComplete(executionAttemptID: ExecutionAttemptID, jobID: JobID, checkpointID: Long, timestamp: Long, mgr: Remote[TaskManager]) =>
            log.debug(s"Receiver ConfirmedCheckpoint $checkpointID@$timestamp " +
                s"for $executionAttemptID")
            val task = runningTasks.get(executionAttemptID)
            if (task != null) {
                task.notifyCheckpointComplete(checkpointID)
            } else {
                log.debug(s"TaskManager received a checkpoint confirmation " +
                    s"for unknown task $executionAttemptID")
            }
        case TriggerCheckpoint(executionAttemptID: ExecutionAttemptID, jobID: JobID, checkpointID: Long, timestamp: Long, checkpointOptions: CheckpointOptions, mgr: Remote[TaskManager]) =>
            triggerCheckpoint(executionAttemptID, jobID, checkpointID, timestamp, checkpointOptions, mgr)
    }
}
```

```
@multicase trait TaskDistributionSystem {
    type JobManager <: { type Type <: Multiple[TaskManager] }
    type TaskManager <: { type Type <: Single[JobManager] }
}
def disconnectFromJobManager(instanceID: InstanceID, cause: Exception,
    mgr: Remote[TaskManager]): Unit = on[JobManager] {
    case Left(stackTrace) => handleJobManagerDisconnect("JobManager requested disconnect: " + cause)
    case Right(stackTrace) => handleJobManagerDisconnect("JobManager requested disconnect: " + cause)
}
def triggerTaskManagerRegistration(): Unit = {
    triggerTaskManagerRegistration("triggerTaskManagerRegistration")
}
else {
    log.debug(s"Received disconnect message for wrong instance id " + instanceID)
}

def stopCluster(applicationStatus: ApplicationStatus, message: String,
    mgr: Remote[TaskManager]): Unit = on[JobManager] {
    case Left(stackTrace) => handleRequestTaskManagerLog(stackTrace, currentJobManager.get)
    case Right(stackTrace) => handleRequestTaskManagerLog(stackTrace, currentJobManager.get)
}
def requestStackTrace(mgr: Remote[TaskManager]): Unit = on[JobManager] {
    case Left(stackTrace) => mgr ! stackTrace
    case Right(stackTrace) => mgr ! stackTrace
}
local.map(_.left.get)

def submitTask(tdd: TaskDeploymentDescriptor): Unit = on[JobManager] {
    case Left(runnableTask) => runnableTask(tdd)
    case Right(stackTrace) => stackTrace
}
local.map(_.left.get)

def stopTask(executionAttemptID: ExecutionAttemptID, mgr: Remote[TaskManager]): Unit = on[JobManager] {
    val task = runningTasks.get(executionAttemptID)
    if (task != null) {
        task.stopExecution()
        Left(Acknowledge.get())
    } else {
        log.debug(s"Cannot find task to stop for execution $executionAttemptID")
    }
    local.map(_.left.get)
}

def cancelTask(executionAttemptID: ExecutionAttemptID, mgr: Remote[TaskManager]): Unit = on[JobManager] {
    val task = runningTasks.get(executionAttemptID)
    if (task != null) {
        task.cancelExecution()
        Left(Acknowledge.get())
    } else {
        log.debug(s"Cannot find task to cancel for execution $executionAttemptID")
    }
    local.map(_.left.get)
}

def updatePartitions(executionAttemptID: ExecutionAttemptID, partitionInfos: java.util.List[PartitionInfo], mgr: Remote[TaskManager]): Unit = on[JobManager] {
    partitionInfos.foreach { info => updateTaskInputPartitions(executionAttemptID, partitionInfos) }
    local.map(_.left.get)
}

def failPartition(executionAttemptID: ExecutionAttemptID, mgr: Remote[TaskManager]): Unit = on[JobManager] {
    log.info(s"Discarding the results produced by task execution $executionAttemptID")
    network.getResultPartitionManager.releasePartitionsProducedBy(executionAttemptID)
    catch {
        case t: Throwable => killTaskManagerFatal(
            s"Fatal leak: Unable to release intermediate result partition data", t)
    }
}

def notifyCheckpointComplete(executionAttemptID: ExecutionAttemptID, jobID: JobID, checkpointID: Long, timestamp: Long, mgr: Remote[TaskManager]): Unit = on[JobManager] {
    log.debug(s"Receiver ConfirmedCheckpoint $checkpointID@$timestamp " +
        s"for $executionAttemptID")
    val task = runningTasks.get(executionAttemptID)
    if (task != null) {
        task.notifyCheckpointComplete(checkpointID)
    } else {
        log.debug(s"TaskManager received a checkpoint confirmation " +
            s"for unknown task $executionAttemptID")
    }
}

def triggerCheckpoint(executionAttemptID: ExecutionAttemptID, jobID: JobID, checkpointID: Long, timestamp: Long, checkpointOptions: CheckpointOptions, mgr: Remote[TaskManager]): Unit = on[JobManager] {
    triggerCheckpoint(executionAttemptID, jobID, checkpointID, timestamp, checkpointOptions, mgr)
}
```



Crosscutting functionality separated among compilation units

```

class TaskManagerGateway {
    def handleFromJobManager(instanceId: InstanceID, cause: Exception,
                           m: ? JobManagerRef)(instanceId, cause) {
        m ! JobManagerRef(instanceId, cause)
    }
    def handleCluster(applicationStatus: ApplicationStatus, message: String,
                      m: ? JobManagerRef)(applicationStatus, message) {
        m ! JobManagerRef(applicationStatus, message)
    }
    def handleTask(tdd: TaskDeploymentDescriptor, mgr: ActorRef) = {
        (m: ? JobManagerRef).mapTo[Acknowledge]
    }
    def handleTaskAttempt(tdd: TaskDeploymentDescriptor, mgr: ActorRef) = {
        (m: ? JobManagerRef).mapTo[Acknowledge]
    }
    def handleTaskExecutionAttemptID(executionAttemptID: ExecutionAttemptID, mgr: ActorRef) = {
        (m: ? JobManagerRef).mapTo[Acknowledge]
    }
    def handleTaskExecutionAttemptID(partitionInfo: PartitionInfo, mgr: ActorRef) = {
        (m: ? JobManagerRef).mapTo[Acknowledge]
    }
    def handleTaskExecutionAttemptID(jobID: JobID, executionAttemptID: ExecutionAttemptID, mgr: ActorRef) = {
        (m: ? JobManagerRef).mapTo[Acknowledge]
    }
    def handleTaskExecutionAttemptID(jobID: JobID, executionAttemptID: ExecutionAttemptID, checkpointOptions: CheckpointOptions, mgr: ActorRef) = {
        (m: ? JobManagerRef).mapTo[Acknowledge]
    }
    def handleTaskExecutionAttemptID(jobID: JobID, executionAttemptID: ExecutionAttemptID, timestamp: Long, checkpointOptions: CheckpointOptions, mgr: ActorRef) = {
        (m: ? JobManagerRef).mapTo[Acknowledge]
    }
    def handleTaskExecutionAttemptID(jobID: JobID, executionAttemptID: ExecutionAttemptID, timestamp: Long, checkpointOptions: CheckpointOptions, m: ? JobManagerRef)(jobID, executionAttemptID, timestamp, checkpointOptions, mgr: ActorRef) = {
        (m: ? JobManagerRef).mapTo[Acknowledge]
    }
}

class TaskManager extends Actor {
    def receive = {
        case StackTraceXTrace => sendToCorrelator(StackTraceXTrace)
        case StopCluster(<*>) => sendToCorrelator(StopCluster)
        case Disconnect(instanceIdToDisconnect, cause) =>
            if (instanceIdToDisconnect.equals(instanceId)) {
                handleJobManagerDisconnect("JobManager requested disconnect: " + triggerJobManagerRegistration())
            } else {
                log.debug("Received disconnect message for wrong instance id " + instanceIdToDisconnect)
            }
        case StopCluster(applicationStatus, message) =>
            log.info(s"Stopping taskmanager with final application status ${applicationStatus} and diagnostics: $message")
            shutdown()
        case RequestTaskManagerLog(requestType) =>
            blobService match {
                case Left(message) => sender ! message
                case Right(message) => sender ! message
            }
        case NoAvailable(<*>) => sender ! akka.actor.Status.Failure(new IOException("Blob storage not available. Cannot upload TaskManager logs."))
        case SubmitTask(tdd: TaskDeploymentDescriptor, mgr: ActorRef) =>
            val task = runningTasks.get(taskExecutionId)
            if (task != null) {
                task.cancelExecution()
                task.correlateMessage(Acknowledge.get())
            } else {
                task ! akka.actor.Status.Failure(t)
            }
        case CancelTask(executionID) =>
            if (runningTasks.get(executionID) != null) {
                runningTasks.get(executionID).cancelExecution()
                runningTasks.get(executionID).correlateMessage(Acknowledge.get())
            } else {
                log.debug(s"Cannot find task to cancel for execution $executionID")
                sender ! akka.actor.Status.Failure(t)
            }
        case TriggerCheckpoint(jobID: JobID, taskExecutionID, checkpointID, timestamp, checkpointOptions) =>
            log.debug(s"Triggering checkpoint $checkpointID@$timestamp for unknown task $taskExecutionID")
            val task = runningTasks.get(taskExecutionID)
            if (task != null) {
                task.triggerCheckpointBarrier(checkpointID, timestamp, checkpointOptions)
            } else {
                log.debug(s"TaskManager received a checkpoint request " + s"for unknown task $taskExecutionID")
            }
        case NotifyCheckpointComplete(jobID, taskExecutionID, checkpointID, timestamp) =>
            log.debug(s"Receiver ConfirmsCheckpoint $checkpointID@$timestamp for $taskExecutionID")
            val task = runningTasks.get(taskExecutionID)
            if (task != null) {
                task.notifyCheckpointComplete(checkpointID)
            } else {
                log.debug(s"TaskManager received a checkpoint confirmation " + s"for unknown task $taskExecutionID")
            }
    }
}

```



Flink



Developers are **not** forced to modularize **along network boundaries**

```

@multicard trait TaskDistributionSystem {
    type JobManager <: { type Type <: Multiple[TaskManager] }
    type TaskManager <: { type Type <: Single[JobManager] }
    def handleFromJobManager(instanceId: InstanceID, cause: Exception,
                           m: ? JobManagerRef)(instanceId, cause) {
        m ! JobManagerRef(instanceId, cause)
    }
    def handleJobManagerDisconnect(<*>) => handleJobManagerDisconnect("JobManager requested disconnect: " + triggerJobManagerRegistration())
    else {
        log.debug("Received disconnect message for wrong instance id " + instanceId)
    }
}

def stopCluster(applicationStatus: ApplicationStatus, message: String,
               m: ? JobManagerRef)(applicationStatus, message) =>
    log.info(s"Stopping taskmanager with final application status ${applicationStatus} and diagnostics: $message")
    shutdown()

def stopTask(executionAttemptID: ExecutionAttemptID, m: ? JobManagerRef) =>
    m ! JobManagerRef(executionAttemptID)
    local.map_.left.get()

def submitTask(tdd: TaskDeploymentDescriptor, m: ? JobManagerRef) =>
    m ! JobManagerRef(tdd)
    local.map_.left.get()

def stopTask(executionAttemptID: ExecutionAttemptID, m: ? JobManagerRef) =>
    m ! JobManagerRef(executionAttemptID)
    local.map_.left.get()

```



That's only half the battle!

How to modularize code into (distributed) system functionalities?

```

def cancelTask(executionID) =>
    if (runningTasks.get(executionID) != null) {
        runningTasks.get(executionID).cancelExecution()
        runningTasks.get(executionID).correlateMessage(Acknowledge.get())
    } else {
        log.debug(s"Cannot find task to cancel for execution $executionID")
        sender ! akka.actor.Status.Failure(t)
    }

def failPartition(executionAttemptID: ExecutionAttemptID, jobID: JobID, timestamp: Long, m: ? JobManagerRef)(executionAttemptID, jobID, timestamp) =>
    log.info(s"Discarding the results produced by task execution $executionID")
    network.getResultPartitionManager.releasePartitionsProducedBy(executionID)
    catch {
        case t: Throwable => killTaskManagerFatal(
            "Fatal leak: Unable to release intermediate result partition data", t)
    }

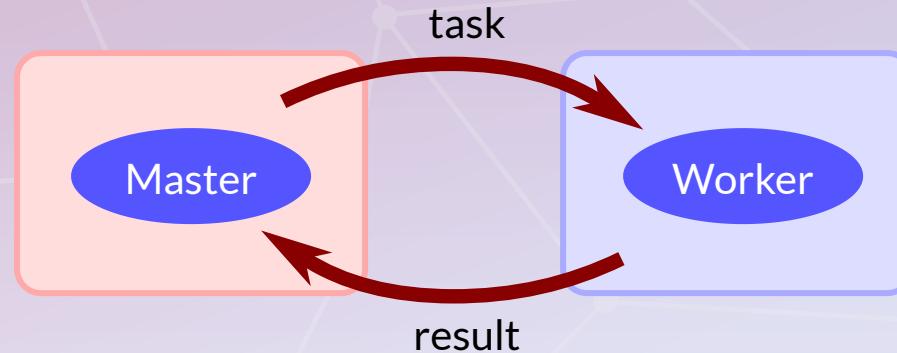
def notifyCheckpointComplete(executionAttemptID: ExecutionAttemptID, jobID: JobID, checkpointID: Long, timestamp: Long, m: ? JobManagerRef)(executionAttemptID, jobID, checkpointID, timestamp) =>
    task = runningTasks.get(executionAttemptID)
    if (task != null) {
        task.notifyCheckpointComplete(checkpointID)
    } else {
        log.debug(s"TaskManager received a checkpoint confirmation " + s"for unknown task $taskExecutionID")
    }

def triggerCheckpoint(executionAttemptID: ExecutionAttemptID, jobID: JobID, timestamp: Long, checkpointOptions: CheckpointOptions, m: ? JobManagerRef)(executionAttemptID, jobID, timestamp, checkpointOptions) =>
    log.debug(s"Triggering checkpoint $checkpointID@$timestamp for executionAttemptID")
    task = runningTasks.get(executionAttemptID)
    if (task != null) {
        task.triggerCheckpointBarrier(timestamp)
    } else {
        log.debug(s"TaskManager received a checkpoint request " + s"for executionAttemptID")
    }

```

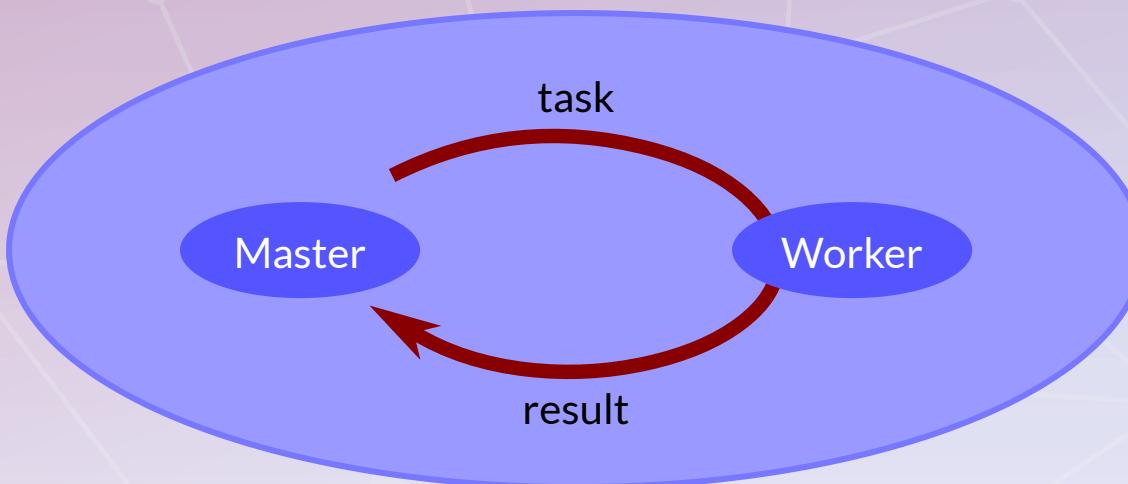
# Distributed Functionalities

---



# Distributed Multitier Functionalities

---



# Multitier Modules

---

- Handle large code bases
- Modularize distributed system functionalities
- Compose subsystems

Abstract Peer Types

disentangle  
distribution  
and  
modularization

# Abstract Peer Types

---

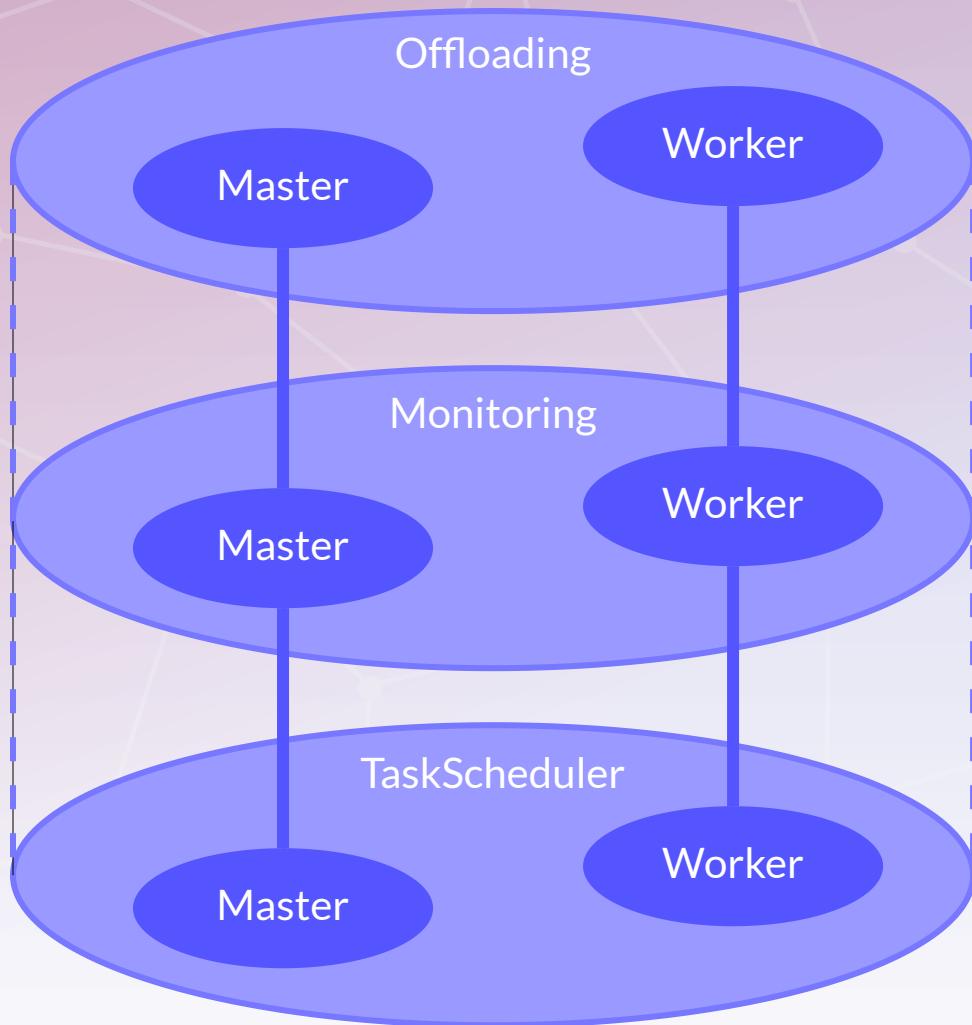
- Define multitier modules on abstract peer types
- Compose functionality of different modules by identifying abstract peer types

Abstract Peer Types

Scala Traits and Objects

disentangle  
distribution  
and  
modularization

# Stacking Multitier Modules



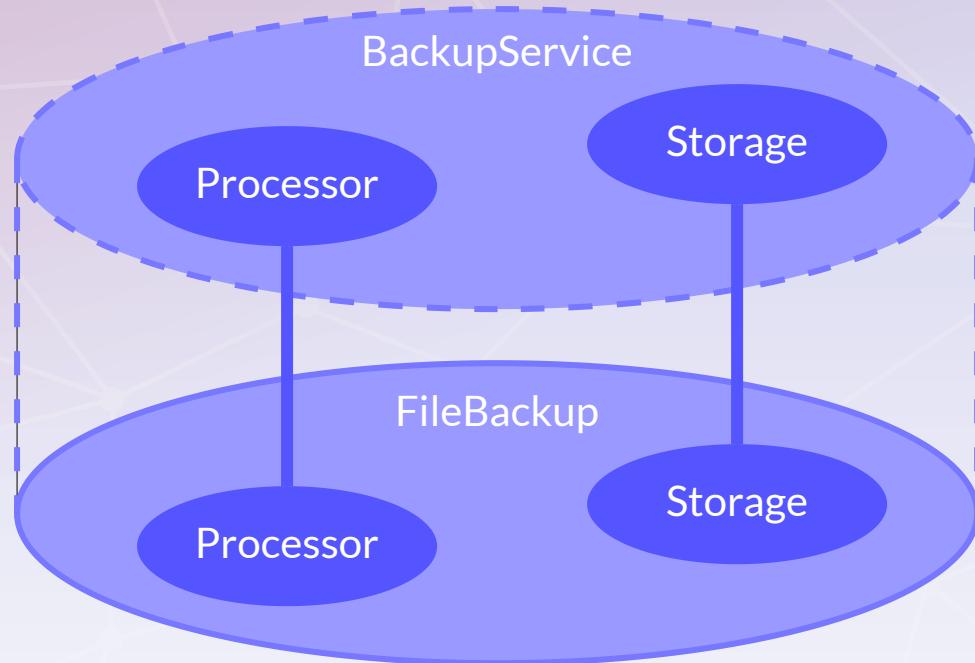
```
@multitier trait Offloading[T] {
    @peer type Master <: { type Tie <: Multiple[Worker] }
    @peer type Worker <: { type Tie <: Single[Master] }
    def run(task: Task[T]): Future[T] on Master =
        placed { (remote(selectWorker()) call execute(task)).asLocal }
    private def execute(task: Task[T]): T on Worker =
        placed { task.process() }
}

@multitier trait Monitoring {
    @peer type Master <: { type Tie <: Multiple[Worker] }
    @peer type Worker <: { type Tie <: Single[Master] }
    def monitoredTimedOut(monitored: Remote[Worker]): Unit on Master
}

@multitier trait TaskScheduler[T] extends
    Offloading[T] with
    Monitoring
```

# Abstract Multitier Modules

---



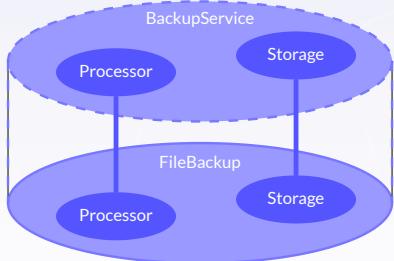
# Interfaces for Subsystems

```
@multitier trait BackupService {  
    @peer type Processor <: { type Tie <: Single[Storage] }  
    @peer type Storage <: { type Tie <: Single[Processor] }  
  
    upper bound allows for refinement  
  
    def store(id: Long, data: Data): Unit on Processor  
    def load(id: Long): Future[Data] on Processor  
}
```

Subsystem  
Architecture

Placed Methods

Modularization  
Across Peers



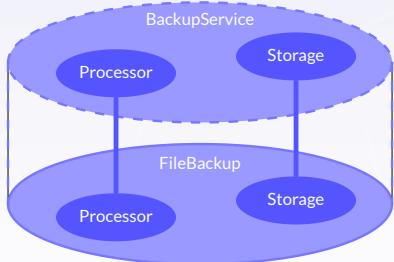
# Implementations for Subsystems

```
@multitier trait FileBackup extends BackupService {  
    def store(id: Long, data: Data): Unit on Processor =  
        placed { remote call write(id, data) }  
    def load(id: Long): Future[Data] on Processor =  
        placed { (remote call read(id)).asLocal }
```

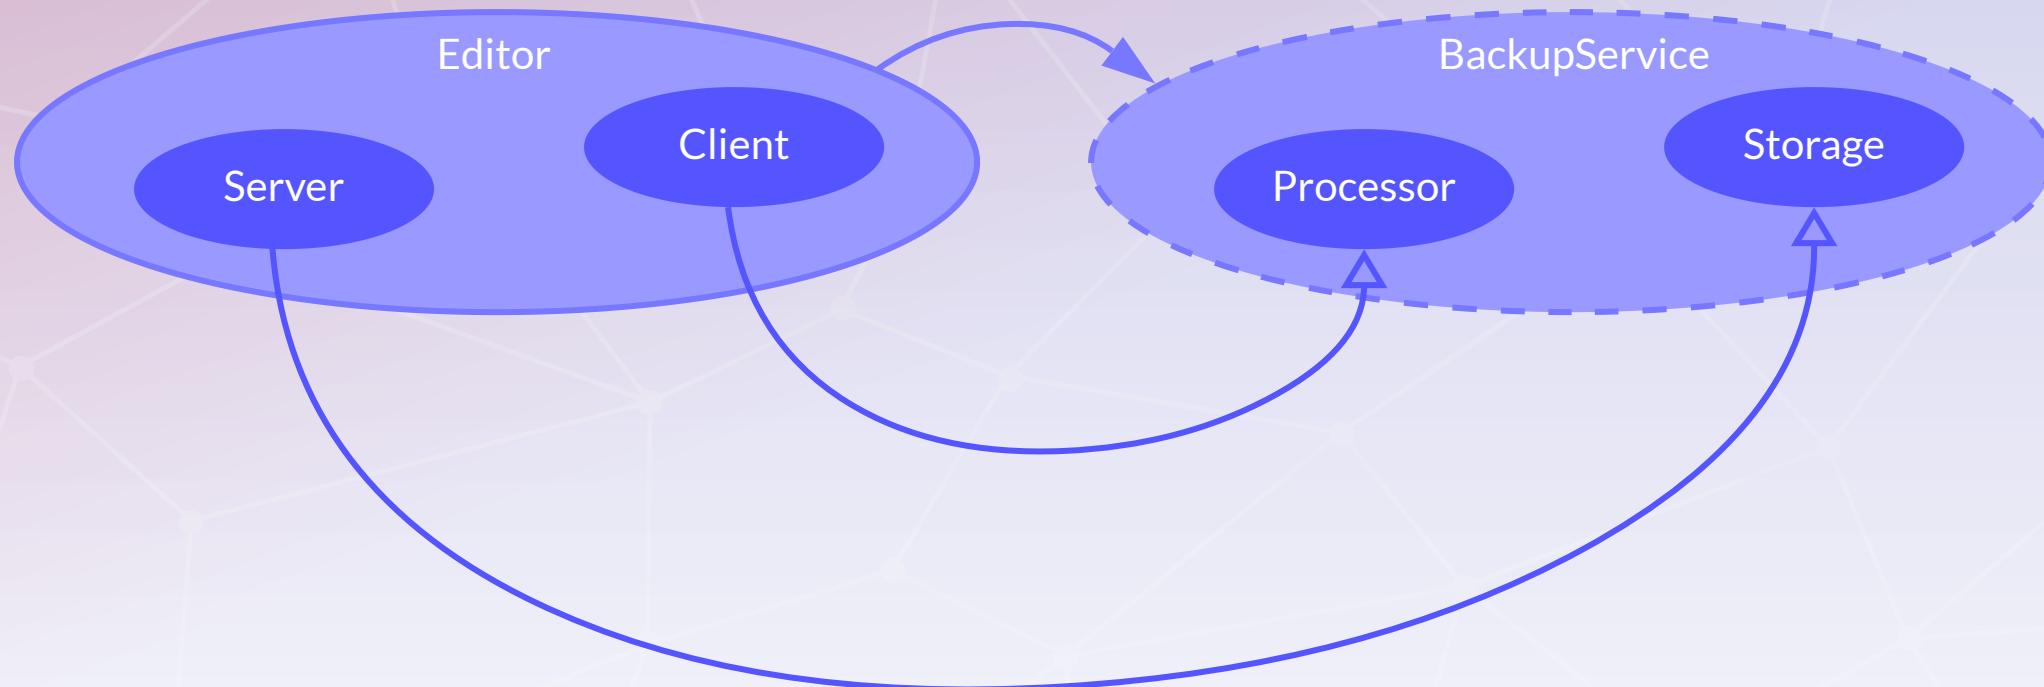
Implementation for  
Abstract Methods

```
private def write(id: Long, data: Data): Unit on Storage =  
    placed { writeToFile(data, s"/storage/$id") }  
private def read(id: Long): Data on Storage =  
    placed { readFromFile[Data](s"/storage/$id") }  
}
```

Encapsulation

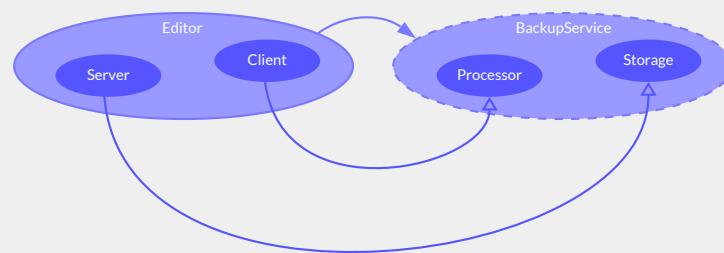


# References to Multitier Modules



# Composing Multitier Modules by References

```
@multitier trait Editor {  
    abstract module reference  
    val backup: BackupService  
    peer refinement  
    @peer type Client <: backup.Processor {  
        type Tie <: Single[Server] with Single[backup.Storage] }  
    @peer type Server <: backup.Storage {  
        type Tie <: Single[Client] with Single[backup.Processor] }  
    }  
  
    @multitier object editor extends Editor {  
        @multitier object backup extends FileBackup  
    }
```



Module Composition

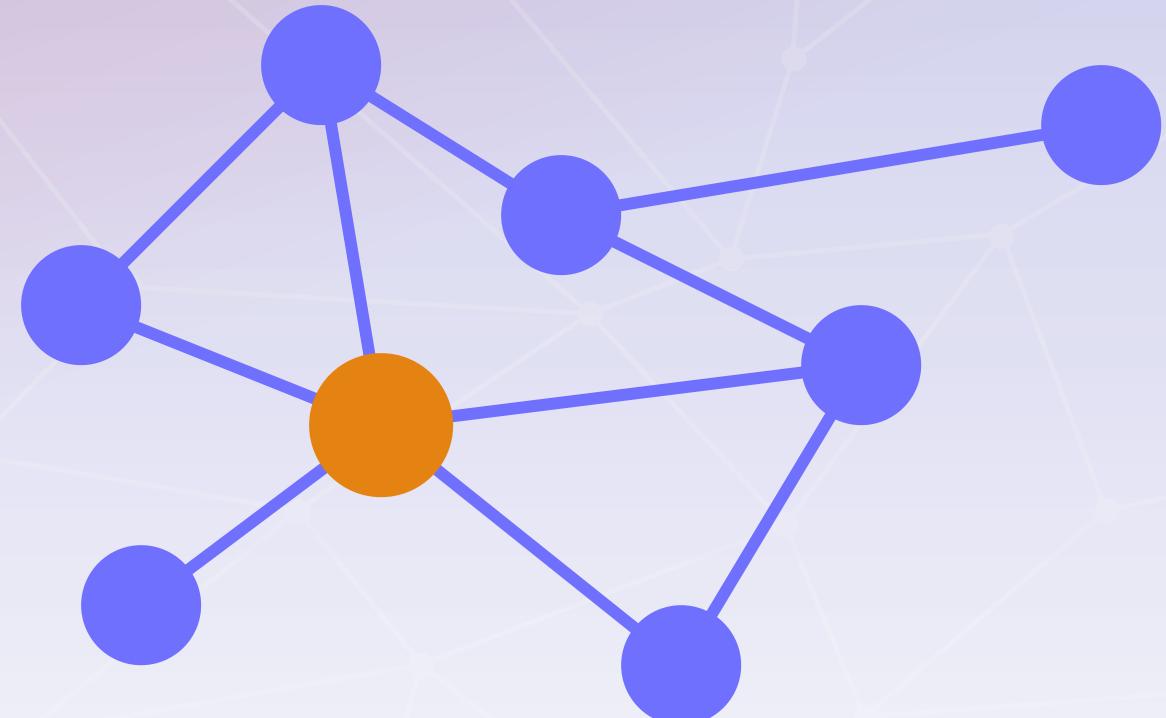
Peer Composition

Instantiation

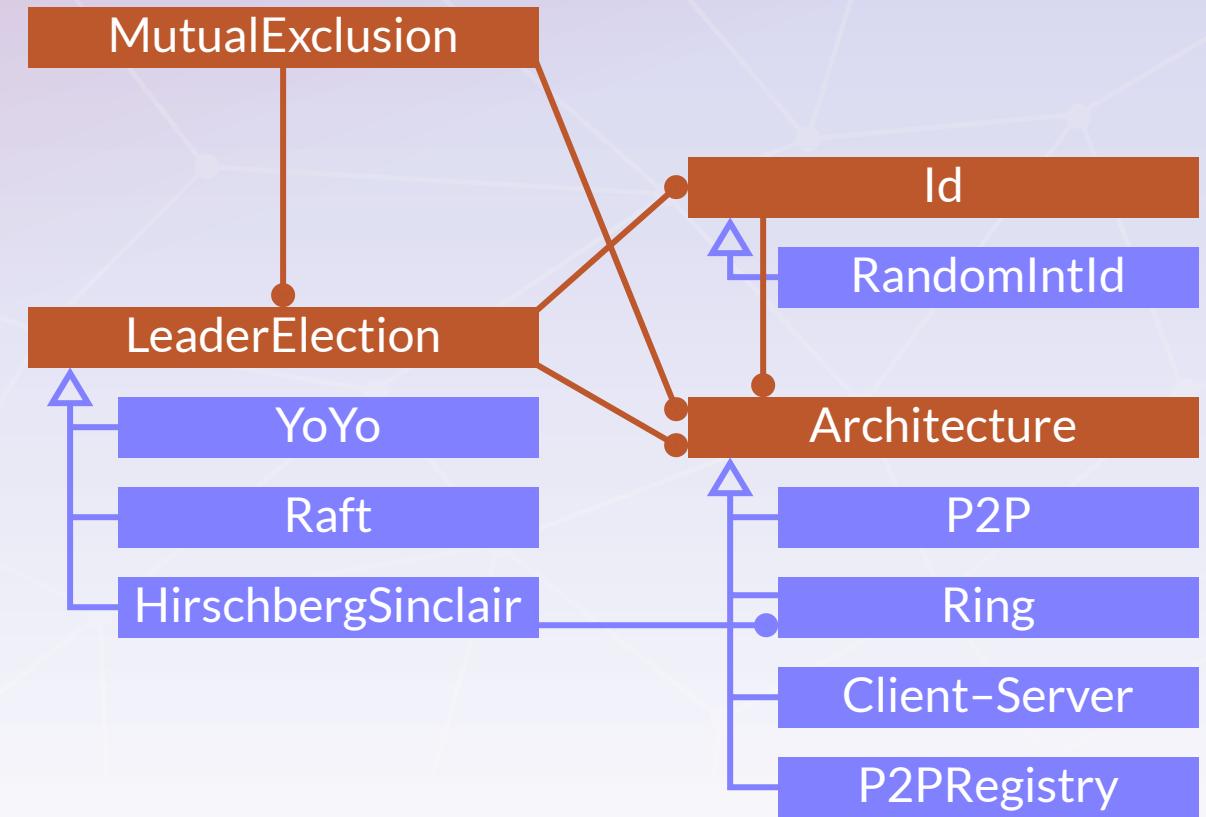
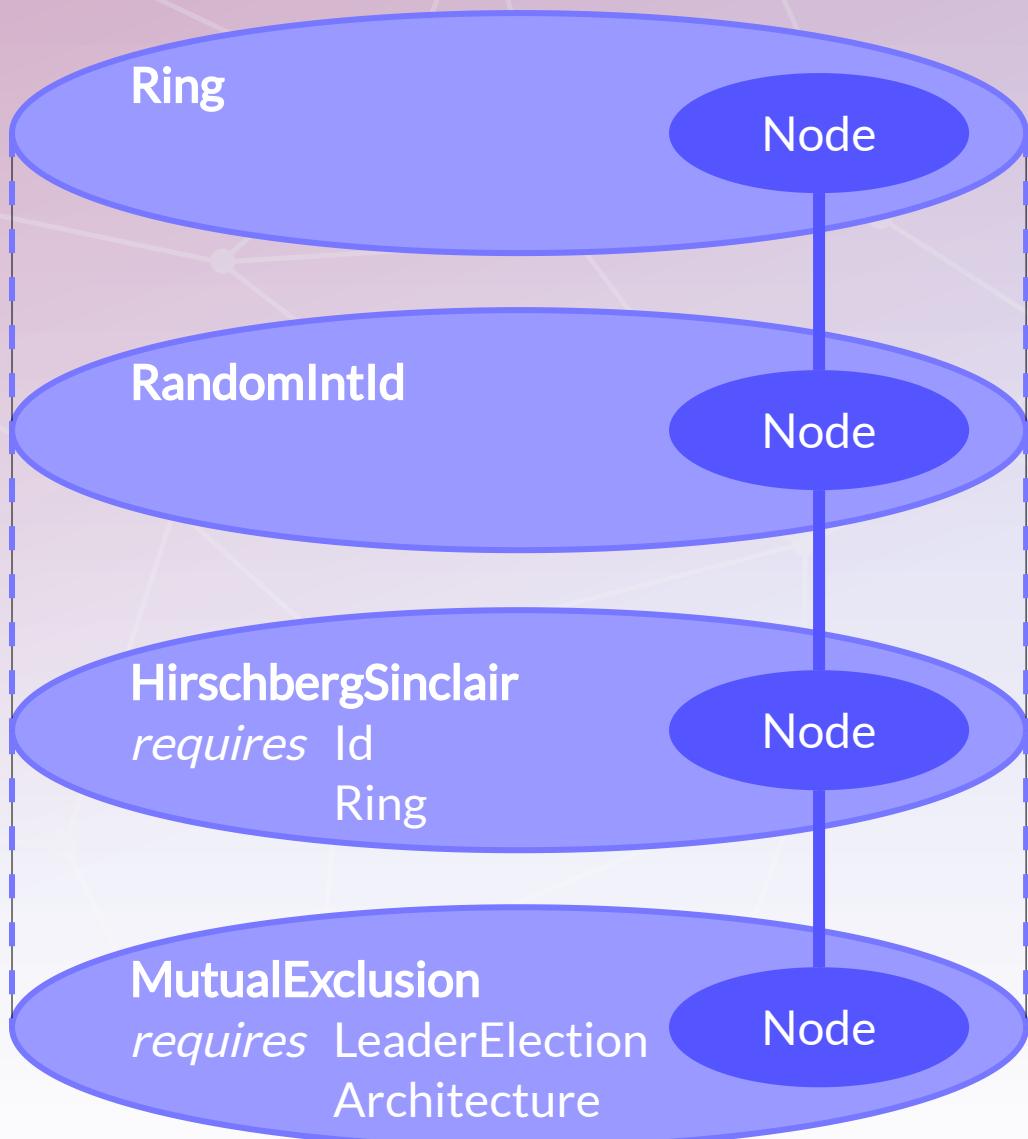
# Distributed Mutual Exclusion Algorithm

---

- Nodes elect a leader
- Followers acquire locks
- Leader grants or denies the lock



# Mixing Constrained Modules



# Leader Election Case Study

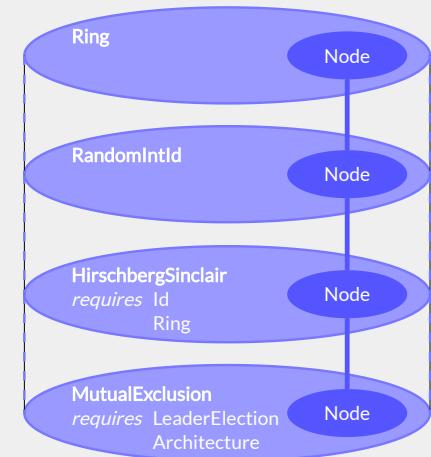
```
@multitier trait MutualExclusion[T] {
    this: Architecture with LeaderElection[T] =>
    def lock(id: T): Boolean on Node = { ... }
    def unlock(id: Id): Unit on Node = { ... }
}

@multitier trait LeaderElection[T] {
    this: Architecture with Id[T] =>
    def electLeader(): Unit on Node
    def electedAsLeader(): Unit on Node
}

@multitier abstract class Id[T: Ordering] {
    this: Architecture =>
    val id: Local[T] on Node
}
```

```
@multitier trait HirschbergSinclair[T]
extends LeaderElection[T] {
    this: Ring with Id[T] =>
    def electLeader() = on[Node] { elect(0) }
    private def elect(phase: Int) = on[Node] { /* ... */ }
    private def propagate(remoteId: T, hops: Int,
        direction: Direction) = on[Node] { /* ... */ }
}

@multitier object locking extends
    MutualExclusion[Int] with
    HirschbergSinclair[Int] with
    Ring with
    RandomIntId
```



```

@multitier trait TaskDistributionSystem extends CheckpointResponder with KvStateRegistryListener with PartitionProducerStateChecker with ResultPartitionConsumableNotifier with TaskManagerGateway with TaskManagerAction

@multitier trait TaskManagerActions {
  @peer type JobManager <: { type Tie <: Multiple[TaskManager] }
  @peer type TaskManager <: { type Tie <: Single[JobManager] }

  notifyState(executionAttemptID: ExecutionAttemptID) => {
    [taskManager].run.capture(executionAttemptID) {
      unregisterTaskAndNotifyFinalState(executionAttemptID)
    }
  }

  interceptDisconnect(message: String, cause: Throwable) => on[TaskManager] {
    [taskManager].run.capture(message, cause) {
      if (instanceId.equals(instanceID)) {
        [taskManager].run.capture(disconnect, cause) {
          triggerTaskManagerRegistration()
        }
      } else {
        log.debug("Received disconnect message for wrong instance id " + instanceID)
      }
    }
  }

  def stopCluster(applicationStatus: ApplicationStatus, message: String, mgr: Remote[TaskManager]) = on[JobManager] {
    [mgr].run.capture(tdd) {
      log.info(s"Stopping JobManager with final application status " + applicationStatus and diagnostics: message)
      shutdown()
    }
  }

  def cancelTask(mgr: Remote[TaskManager]) = on[JobManager] {
    [mgr].run.capture(tdd) {
      [mgr].cancelTask(mgr)
    }
  }

  def submitTask(tdd: TaskDeploymentDescriptor, mgr: Remote[TaskManager]) = on[JobManager] {
    [mgr].run.capture(tdd) {
      [mgr].submitTask(mgr)
    }
  }

  def updateTaskExecutionState(taskExecutionState: TaskExecutionState) = on[TaskManager] {
    [taskManager].run.capture(taskExecutionState) {
      currentJobManager.foreach { jobManager =>
        val future = jobManager.respondToUpdateTaskExecutionState(taskExecutionState)(askTimeout)
        future.map[Boolean].onComplete {
          scala.util.Success(result) =>
            if (!result) {
              log.error("Failed task execution message")
            } else {
              taskExecutionState.getID,
              new Exception("Task has been cancelled on the JobManager.")
            }
        }
        case scala.util.Failure(t) =>
          self ! decorateMessage(FailTask(
            task = runningTasks.get(executionAttemptID),
            exception = Exception("Failed to send ExecutionStateChange notification " +
              "to JobManager", t)))
        case t: Throwable =>
          Right(Status.Failure(t))
      }
    }
  }

  def cancelTask(executionAttemptID: ExecutionAttemptID, mgr: Remote[TaskManager]) = on[JobManager] {
    [mgr].run.capture(executionAttemptID) {
      task = runningTasks.get(executionAttemptID)
      if (task != null) {
        task.stopExecution()
        Left(Acknowledge.get())
      } else {
        log.debug("Cannot find task to stop for execution $executionAttemptID")
        Left(Acknowledge.get())
      }
    }
  }

  def updatePartitions(executionAttemptID: ExecutionAttemptID, partitionInfos: java.util.List[PartitionInfo], mgr: Remote[TaskManager]) = on[JobManager] {
    [mgr].run.capture(executionAttemptID, partitionInfos) {
      currentJobs.get(jobID).map(_.left.get)
    }
  }

  def failPartition(executionAttemptID: ExecutionAttemptID, mgr: Remote[TaskManager]) = on[JobManager] {
    [mgr].run.capture(executionAttemptID) {
      log.info("Discarding the results produced by task execution $executionID")
      network.getResultPartitionManager.releasePartitionsProducedBy(executionID)
    }
    catch {
      case t: Throwable => killTaskManagerFatal(
        s"Fatal leak: Unable to release intermediate result partition data.", t)
    }
  }

  def notifyCheckpointComplete(executionAttemptID: ExecutionAttemptID, jobID: JobID, timestamp: Long, checkpointID: Long, checkpointOptions: CheckpointOptions, mgr: Remote[TaskManager]) = on[JobManager] {
    [mgr].run.capture(executionAttemptID, jobID, checkpointID, timestamp) {
      log.debug("TaskManager received a checkpoint confirmation " +
        s"for unknown task $taskExecutionID")
    }
  }

  def triggerCheckpoint(executionAttemptID: ExecutionAttemptID, jobID: JobID, checkpointID: Long, timestamp: Long, checkpointOptions: CheckpointOptions, mgr: Remote[TaskManager]) = on[JobManager] {
    [mgr].run.capture(executionAttemptID, jobID, checkpointID, timestamp, checkpointOptions) {
      log.debug("TaskManager received a triggerCheckpoint $checkpointID@$timestamp " +
        s"for executionAttemptID")
    }
  }

  @multitier trait TaskManagerGateway {
    @peer type JobManager <: { type Tie <: Multiple[TaskManager] }
    @peer type TaskManager <: { type Tie <: Single[JobManager] }

    if (jobManager.isConnected(instanceID: JobID, partitionID: ResultPartitionID) && jobManager.isAlive) {
      jobManager.run.capture(instanceID) {
        jobManager.capture(instanceID, cause: Exception, graph: Graph)
      }
    }
  }

  interceptDisconnect(message: String, cause: Throwable) = on[TaskManager] {
    [taskManager].run.capture(message, cause) {
      if (jobManager.isConnected(instanceID: JobID, partitionID: ResultPartitionID) && jobManager.isAlive) {
        jobManager.requestedDisconnect := true
        jobManager.run.capture(instanceID, cause: Exception, graph: Graph)
      }
    }
  }

  failTask(executionAttemptID: ExecutionAttemptID, cause: Throwable) = on[TaskManager] {
    [taskManager].run.capture(executionAttemptID, cause) {
      if (jobManager.isConnected(instanceID: JobID, partitionID: ResultPartitionID) && jobManager.isAlive) {
        jobManager.requestedDisconnect := true
        jobManager.run.capture(instanceID, cause: Exception, graph: Graph)
      }
    }
  }

  stopCluster(applicationStatus: ApplicationStatus, message: String, mgr: Remote[TaskManager]) = on[JobManager] {
    [mgr].run.capture(tdd) {
      log.info(s"Stopping JobManager with final application status " + applicationStatus and diagnostics: message)
      shutdown()
    }
  }

  updateTaskExecutionState(taskExecutionState: TaskExecutionState) = on[TaskManager] {
    [taskManager].run.capture(taskExecutionState) {
      currentJobManager.foreach { jobManager =>
        val future = jobManager.respondToUpdateTaskExecutionState(taskExecutionState)(askTimeout)
        future.map[Boolean].onComplete {
          scala.util.Success(result) =>
            if (!result) {
              log.error("Failed task execution message")
            } else {
              taskExecutionState.getID,
              new Exception("Task has been cancelled on the JobManager.")
            }
        }
        case scala.util.Failure(t) =>
          self ! decorateMessage(FailTask(
            task = runningTasks.get(executionAttemptID),
            exception = Exception("Failed to send ExecutionStateChange notification " +
              "to JobManager", t)))
        case t: Throwable =>
          Right(Status.Failure(t))
      }
    }
  }

  cancelTask(executionAttemptID: ExecutionAttemptID, mgr: Remote[TaskManager]) = on[JobManager] {
    [mgr].run.capture(executionAttemptID) {
      task = runningTasks.get(executionAttemptID)
      if (task != null) {
        task.stopExecution()
        Left(Acknowledge.get())
      } else {
        log.debug("Cannot find task to stop for execution $executionAttemptID")
        Left(Acknowledge.get())
      }
    }
  }

  updatePartitions(executionAttemptID: ExecutionAttemptID, partitionInfos: java.util.List[PartitionInfo], mgr: Remote[TaskManager]) = on[JobManager] {
    [mgr].run.capture(executionAttemptID, partitionInfos) {
      currentJobs.get(jobID).map(_.left.get)
    }
  }

  failPartition(executionAttemptID: ExecutionAttemptID, mgr: Remote[TaskManager]) = on[JobManager] {
    [mgr].run.capture(executionAttemptID) {
      log.info("Discarding the results produced by task execution $executionID")
      network.getResultPartitionManager.releasePartitionsProducedBy(executionID)
    }
    catch {
      case t: Throwable => killTaskManagerFatal(
        s"Fatal leak: Unable to release intermediate result partition data.", t)
    }
  }

  notifyCheckpointComplete(executionAttemptID: ExecutionAttemptID, jobID: JobID, timestamp: Long, checkpointID: Long, checkpointOptions: CheckpointOptions, mgr: Remote[TaskManager]) = on[JobManager] {
    [mgr].run.capture(executionAttemptID, jobID, checkpointID, timestamp) {
      log.debug("TaskManager received a checkpoint confirmation " +
        s"for unknown task $taskExecutionID")
    }
  }

  triggerCheckpoint(executionAttemptID: ExecutionAttemptID, jobID: JobID, checkpointID: Long, timestamp: Long, checkpointOptions: CheckpointOptions, mgr: Remote[TaskManager]) = on[JobManager] {
    [mgr].run.capture(executionAttemptID, jobID, checkpointID, timestamp, checkpointOptions) {
      log.debug("TaskManager received a triggerCheckpoint $checkpointID@$timestamp " +
        s"for executionAttemptID")
    }
  }

  requestTaskManager(logTypeRequest: LogTypeRequest, mgr: Remote[TaskManager]) = on[JobManager] {
    [mgr].run.capture(logTypeRequest) {
      obService match {
        case Some(service) =>
          handleRequestTaskManager(logTypeRequest, currentJobManager.get)
        case None =>
          Right(Status.ActorStatusFailure(new IOException("JobManager not available. Cannot upload TaskManager logs.")))
      }
    }
  }

  @multitier trait KvStateRegistryListener {
    @peer type JobManager <: { type Tie <: Multiple[TaskManager] }
    @peer type TaskManager <: { type Tie <: Single[JobManager] }

    if (jobManager.isConnected(jobID: JobID, partitionID: ResultPartitionID) && jobManager.isAlive) {
      jobManager.run.capture(instanceID) {
        kvStateId = kvStateId
        instanceID = instanceID
        keyGroupRange = keyGroupRange
        registrationName = registrationName
        kvStateId, kvStateServerAddress = kvStateId
        currentJobs.get(jobID).match {
          case Some(graph) =>
            try {
              graph.getKVStateRegistered(jobID, instanceID, keyGroupRange, registrationName)
            } catch {
              case e: ExecutionGraph.ScheduleUpdateConsumersException =>
                log.error("Failed to notify KvStateRegistry about registration.")
            }
          case None =>
            log.error("Received state registration for unavailable job.")
        }
      }
    }
  }

  interceptDisconnect(message: String, cause: Throwable) = on[TaskManager] {
    [taskManager].run.capture(message, cause) {
      if (jobManager.isConnected(jobID: JobID, partitionID: ResultPartitionID) && jobManager.isAlive) {
        jobManager.requestedDisconnect := true
        jobManager.run.capture(instanceID, cause: Exception, graph: Graph)
      }
    }
  }

  failTask(executionAttemptID: ExecutionAttemptID, cause: Throwable) = on[TaskManager] {
    [taskManager].run.capture(executionAttemptID, cause) {
      if (jobManager.isConnected(jobID: JobID, partitionID: ResultPartitionID) && jobManager.isAlive) {
        jobManager.requestedDisconnect := true
        jobManager.run.capture(instanceID, cause: Exception, graph: Graph)
      }
    }
  }

  stopCluster(applicationStatus: ApplicationStatus, message: String, mgr: Remote[TaskManager]) = on[JobManager] {
    [mgr].run.capture(tdd) {
      log.info(s"Stopping JobManager with final application status " + applicationStatus and diagnostics: message)
      shutdown()
    }
  }

  updateTaskExecutionState(taskExecutionState: TaskExecutionState) = on[TaskManager] {
    [taskManager].run.capture(taskExecutionState) {
      currentJobManager.foreach { jobManager =>
        val future = jobManager.respondToUpdateTaskExecutionState(taskExecutionState)(askTimeout)
        future.map[Boolean].onComplete {
          scala.util.Success(result) =>
            if (!result) {
              log.error("Failed task execution message")
            } else {
              taskExecutionState.getID,
              new Exception("Task has been cancelled on the JobManager.")
            }
        }
        case scala.util.Failure(t) =>
          self ! decorateMessage(FailTask(
            task = runningTasks.get(executionAttemptID),
            exception = Exception("Failed to send ExecutionStateChange notification " +
              "to JobManager", t)))
        case t: Throwable =>
          Right(Status.Failure(t))
      }
    }
  }

  cancelTask(executionAttemptID: ExecutionAttemptID, mgr: Remote[TaskManager]) = on[JobManager] {
    [mgr].run.capture(executionAttemptID) {
      task = runningTasks.get(executionAttemptID)
      if (task != null) {
        task.stopExecution()
        Left(Acknowledge.get())
      } else {
        log.debug("Cannot find task to stop for execution $executionAttemptID")
        Left(Acknowledge.get())
      }
    }
  }

  updatePartitions(executionAttemptID: ExecutionAttemptID, partitionInfos: java.util.List[PartitionInfo], mgr: Remote[TaskManager]) = on[JobManager] {
    [mgr].run.capture(executionAttemptID, partitionInfos) {
      currentJobs.get(jobID).map(_.left.get)
    }
  }

  failPartition(executionAttemptID: ExecutionAttemptID, mgr: Remote[TaskManager]) = on[JobManager] {
    [mgr].run.capture(executionAttemptID) {
      log.info("Discarding the results produced by task execution $executionID")
      network.getResultPartitionManager.releasePartitionsProducedBy(executionID)
    }
    catch {
      case t: Throwable => killTaskManagerFatal(
        s"Fatal leak: Unable to release intermediate result partition data.", t)
    }
  }

  notifyCheckpointComplete(executionAttemptID: ExecutionAttemptID, jobID: JobID, timestamp: Long, checkpointID: Long, checkpointOptions: CheckpointOptions, mgr: Remote[TaskManager]) = on[JobManager] {
    [mgr].run.capture(executionAttemptID, jobID, checkpointID, timestamp) {
      log.debug("TaskManager received a checkpoint confirmation " +
        s"for unknown task $taskExecutionID")
    }
  }

  triggerCheckpoint(executionAttemptID: ExecutionAttemptID, jobID: JobID, checkpointID: Long, timestamp: Long, checkpointOptions: CheckpointOptions, mgr: Remote[TaskManager]) = on[JobManager] {
    [mgr].run.capture(executionAttemptID, jobID, checkpointID, timestamp, checkpointOptions) {
      log.debug("TaskManager received a triggerCheckpoint $checkpointID@$timestamp " +
        s"for executionAttemptID")
    }
  }

  requestTaskManager(logTypeRequest: LogTypeRequest, mgr: Remote[TaskManager]) = on[JobManager] {
    [mgr].run.capture(logTypeRequest) {
      obService match {
        case Some(service) =>
          handleRequestTaskManager(logTypeRequest, currentJobManager.get)
        case None =>
          Right(Status.ActorStatusFailure(new IOException("JobManager not available. Cannot upload TaskManager logs.")))
      }
    }
  }

  @multitier trait ResultPartitionConsumableNotifier {
    @peer type JobManager <: { type Tie <: Multiple[TaskManager] }
    @peer type TaskManager <: { type Tie <: Single[JobManager] }

    if (jobManager.isConnected(jobID: JobID, partitionID: ResultPartitionID) && jobManager.isAlive) {
      jobManager.run.capture(instanceID) {
        currentJobs.get(jobID).match {
          case Some(graph) =>
            try {
              graph.scheduleUpdateConsumers(partitionID)
            } catch {
              case e: ExecutionGraph.ScheduleUpdateConsumersException =>
                log.error("Failed to schedule or update consumers." + e)
            }
          case None =>
            log.error("Cannot find execution graph for job ID $jobID " +
              "to schedule or update consumers")
        }
      }
    }
  }

  interceptDisconnect(message: String, cause: Throwable) = on[TaskManager] {
    [taskManager].run.capture(message, cause) {
      if (jobManager.isConnected(jobID: JobID, partitionID: ResultPartitionID) && jobManager.isAlive) {
        jobManager.requestedDisconnect := true
        jobManager.run.capture(instanceID, cause: Exception, graph: Graph)
      }
    }
  }

  failTask(executionAttemptID: ExecutionAttemptID, cause: Throwable) = on[TaskManager] {
    [taskManager].run.capture(executionAttemptID, cause) {
      if (jobManager.isConnected(jobID: JobID, partitionID: ResultPartitionID) && jobManager.isAlive) {
        jobManager.requestedDisconnect := true
        jobManager.run.capture(instanceID, cause: Exception, graph: Graph)
      }
    }
  }

  stopCluster(applicationStatus: ApplicationStatus, message: String, mgr: Remote[TaskManager]) = on[JobManager] {
    [mgr].run.capture(tdd) {
      log.info(s"Stopping JobManager with final application status " + applicationStatus and diagnostics: message)
      shutdown()
    }
  }

  updateTaskExecutionState(taskExecutionState: TaskExecutionState) = on[TaskManager] {
    [taskManager].run.capture(taskExecutionState) {
      currentJobManager.foreach { jobManager =>
        val future = jobManager.respondToUpdateTaskExecutionState(taskExecutionState)(askTimeout)
        future.map[Boolean].onComplete {
          scala.util.Success(result) =>
            if (!result) {
              log.error("Failed task execution message")
            } else {
              taskExecutionState.getID,
              new Exception("Task has been cancelled on the JobManager.")
            }
        }
        case scala.util.Failure(t) =>
          self ! decorateMessage(FailTask(
            task = runningTasks.get(executionAttemptID),
            exception = Exception("Failed to send ExecutionStateChange notification " +
              "to JobManager", t)))
        case t: Throwable =>
          Right(Status.Failure(t))
      }
    }
  }

  cancelTask(executionAttemptID: ExecutionAttemptID, mgr: Remote[TaskManager]) = on[JobManager] {
    [mgr].run.capture(executionAttemptID) {
      task = runningTasks.get(executionAttemptID)
      if (task != null) {
        task.stopExecution()
        Left(Acknowledge.get())
      } else {
        log.debug("Cannot find task to stop for execution $executionAttemptID")
        Left(Acknowledge.get())
      }
    }
  }

  updatePartitions(executionAttemptID: ExecutionAttemptID, partitionInfos: java.util.List[PartitionInfo], mgr: Remote[TaskManager]) = on[JobManager] {
    [mgr].run.capture(executionAttemptID, partitionInfos) {
      currentJobs.get(jobID).map(_.left.get)
    }
  }

  failPartition(executionAttemptID: ExecutionAttemptID, mgr: Remote[TaskManager]) = on[JobManager] {
    [mgr].run.capture(executionAttemptID) {
      log.info("Discarding the results produced by task execution $executionID")
      network.getResultPartitionManager.releasePartitionsProducedBy(executionID)
    }
    catch {
      case t: Throwable => killTaskManagerFatal(
        s"Fatal leak: Unable to release intermediate result partition data.", t)
    }
  }

  notifyCheckpointComplete(executionAttemptID: ExecutionAttemptID, jobID: JobID, timestamp: Long, checkpointID: Long, checkpointOptions: CheckpointOptions, mgr: Remote[TaskManager]) = on[JobManager] {
    [mgr].run.capture(executionAttemptID, jobID, checkpointID, timestamp) {
      log.debug("TaskManager received a checkpoint confirmation " +
        s"for unknown task $taskExecutionID")
    }
  }

  triggerCheckpoint(executionAttemptID: ExecutionAttemptID, jobID: JobID, checkpointID: Long, timestamp: Long, checkpointOptions: CheckpointOptions, mgr: Remote[TaskManager]) = on[JobManager] {
    [mgr].run.capture(executionAttemptID, jobID, checkpointID, timestamp, checkpointOptions) {
      log.debug("TaskManager received a triggerCheckpoint $checkpointID@$timestamp " +
        s"for executionAttemptID")
    }
  }

  requestTaskManager(logTypeRequest: LogTypeRequest, mgr: Remote[TaskManager]) = on[JobManager] {
    [mgr].run.capture(logTypeRequest) {
      obService match {
        case Some(service) =>
          handleRequestTaskManager(logTypeRequest, currentJobManager.get)
        case None =>
          Right(Status.ActorStatusFailure(new IOException("JobManager not available. Cannot upload TaskManager logs.")))
      }
    }
  }
}

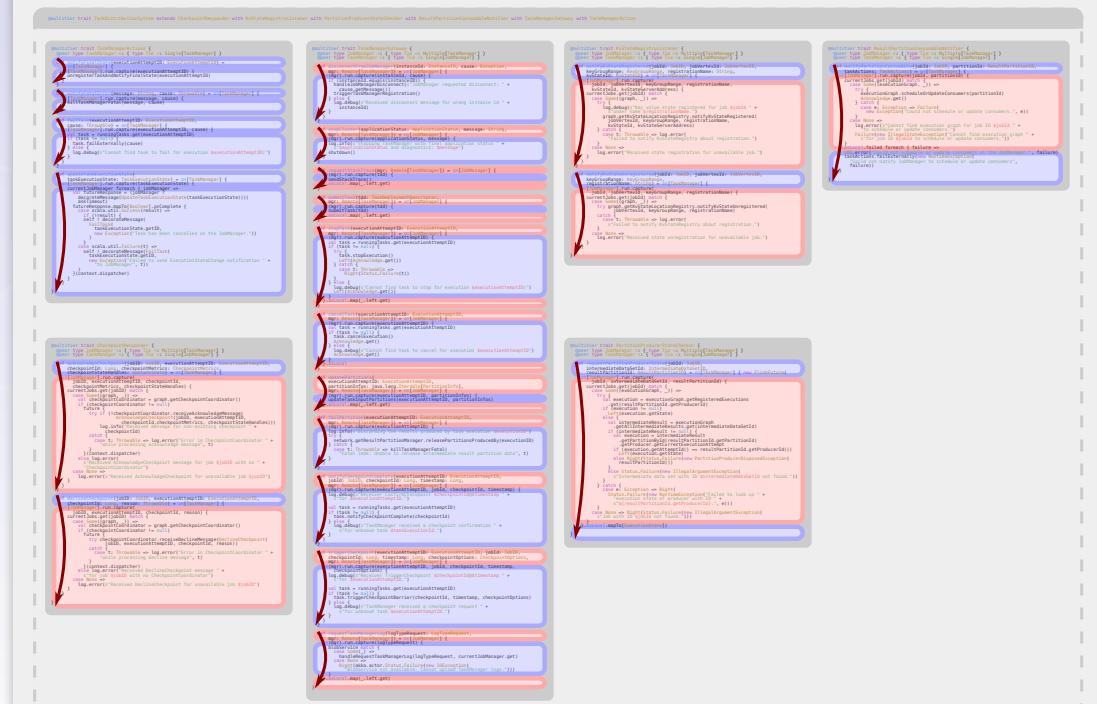


```

# Flink Case Study

```
@multitier object TaskManager {  
    @peer type JobManager <: { type Tie <: Multiple[TaskManager] }  
    @peer type TaskManager <: { type Tie <: Single[JobManager] }  
  
    def submitTask(td: TaskDeployment, tm: Remote[TaskManager]) =  
        on[JobManager] { (remote(tm) call startTask(td)).asLocal }  
    def startTask(td: TaskDeployment) = on[TaskManager] {  
        val task = new Task(td)  
        task.start()  
        Acknowledge()  
    }  
    ...  
  
    @multitier object TaskManagerActions { ... }  
    @multitier object CheckpointResponder { ... }  
    @multitier object ResultPartitionConsumableNotifier { ... }  
    @multitier object PartitionProducerStateChecker { ... }  
    @multitier object KvStateRegistryListener { ... }  
}
```

```
@multitier object TaskDistributionSystem extends  
    TaskManager with  
    TaskManagerActions with  
    CheckpointResponder with  
    ResultPartitionConsumableNotifier with  
    PartitionProducerStateChecker with  
    KvStateRegistryListener
```

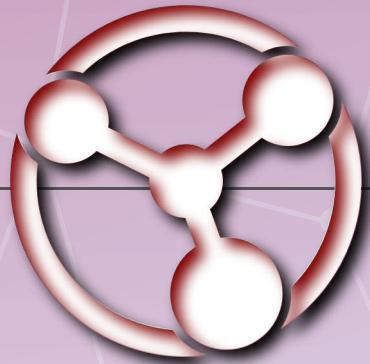


The screenshot shows the homepage of [scala-loci.github.io](https://scala-loci.github.io). The header features a navigation bar with links to Concepts, Step-by-Step Example, Getting Started, Showcases, and Publications. Below the header is a large logo consisting of three spheres connected by a ring, with the text "ScalaLoci" and "A programming language for distributed applications". Three red callout boxes below the logo highlight the language's features: "Unified" (Implement all components of a distributed application in a single language), "Universal" (Freely express any distributed architecture), and "Safe" (Enjoy static type-safety across components and static checks for architectural constraints). The main content area is divided into two sections: "Specify Architecture" (step 1) and "Specify Placement" (step 2). Each section includes a diagram, a brief description, and a code snippet. The "Specify Architecture" section shows a component graph and the following code:

```
@peer type Server <: {  
    type Tie <: Multiple[Client]  
}  
  
@peer type Client <: {  
    type Tie <: Single[Server]  
}
```

The "Specify Placement" section shows a component graph with arrows indicating placement and the following code:

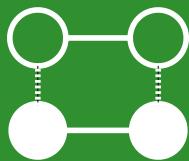
```
val items: Items on Server = placed {  
    getCurrentItems()  
}  
  
val ui: UI on Client = placed {  
    new UI  
}
```



# ScalaLoci: scala-loci.github.io



## Multitier Modules



## Abstract Peer Types

A screenshot of the ScalaLoci website homepage. The header features the text "ScalaLoci" and "A programming language for distributed applications". Below the header are three cards: "Unified" (Implement all components of a distributed application in a single language), "Universal" (Freely express any distributed architecture), and "Safe" (Enjoy static type-safety across components and static checks for architectural constraints). The main content area shows two numbered steps: 1. "Specify Architecture" (Define the architectural relation of the components of the distributed system) with code examples: @peer type Server &lt;: { type Tie &lt;: Multiple[Client] } and @peer type Client &lt;: { type Tie &lt;: Single[Server] }. Step 2. "Specify Placement" (Control where data is located and computations are executed) with code examples: val items: Items on Server = placed { getCurrentItems() } and val ui: UI on Client = placed { new UI() }.

- Pascal Weisenburger, Mirko Köhler, and Guido Salvaneschi. 2018. Distributed System Development with ScalaLoci. *Proceedings of the ACM on Programming Languages* 2, OOPSLA, Article 129.
- Pascal Weisenburger and Guido Salvaneschi. 2019. Multitier Modules. In *Proceedings of the 33rd European Conference on Object-Oriented Programming*, ECOOP.