Hermes, a Fault-Injection Framework (Middleware'13)

Rolando Martins, Rajeev Gandhi, Priya Narasimhan & Soila Pertet (CMU) Antonio Casimiro, Diego Kreutz & Paulo Verissimo (FCUL)

Motivation & Key-Ideas

- **Distributed algorithms hard to implement**
  - Often several parameters need to be tuned but there is no clear insight that helps determine the value of these parameters
- **Fault injectors can help determine problems in implementation and help tune the parameters**
- **Hermes – Our fault-injection framework**
  - **Efficient and lightweight framework**
    - Centralized architecture with lightweight primitives
  - **Ease of use to the developer**
    - Aspect-oriented programming (AOP)
- **Open-source implementation (Java & C++)**
  - **https://github.com/rolandomar/hermes**
  - **https://github.com/rolandomar/hermesCPP**

Possible Applications

- **Paxos based protocols and BFT protocols**
  - Zookeeper, Chubby (Paxos) and BFT-SMaRt (BFT)
- **Chain replication protocols**
  - HDFS block replication, byzantine chain replication
- **Virtual synchrony based group communication**
  - JGroups, Spread, Isis2
- **Distributed file-systems**
  - Ceph, HDFS, PVFS
- **NoSQL**
  - Cassandra, Voldemort ,HBase, MongoDB
- **Orchestration**
  - Mesos, Omega, Openstack Controller
- Software-defined networking (SDN)
  - Neutron+OVS

Architecture

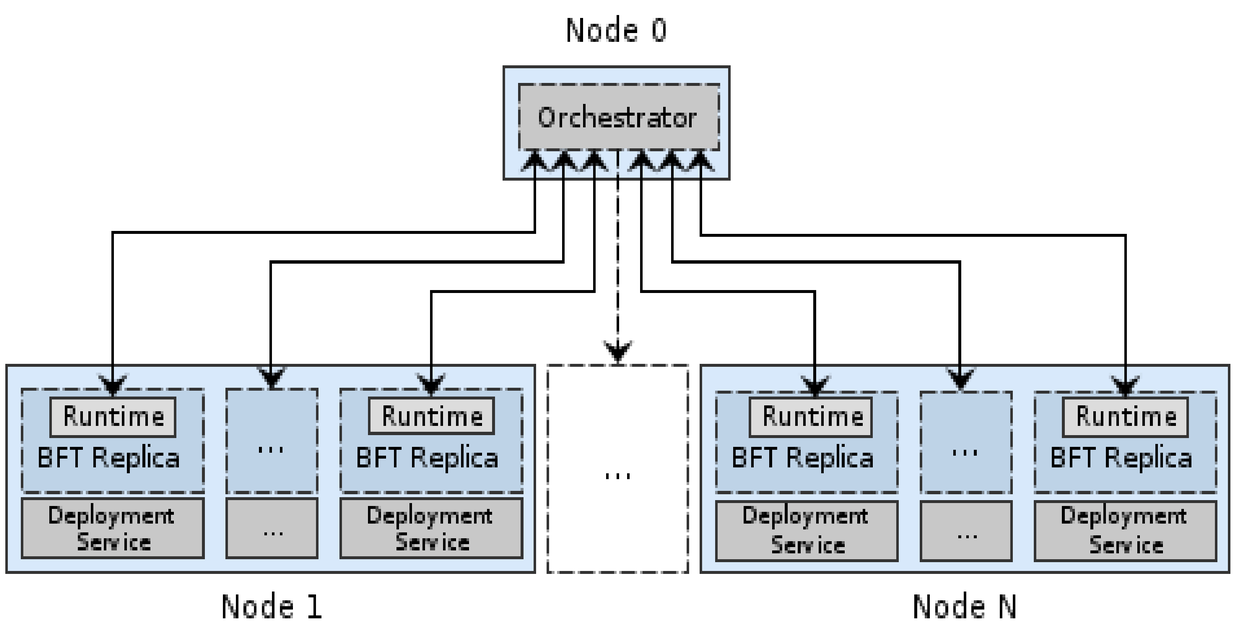

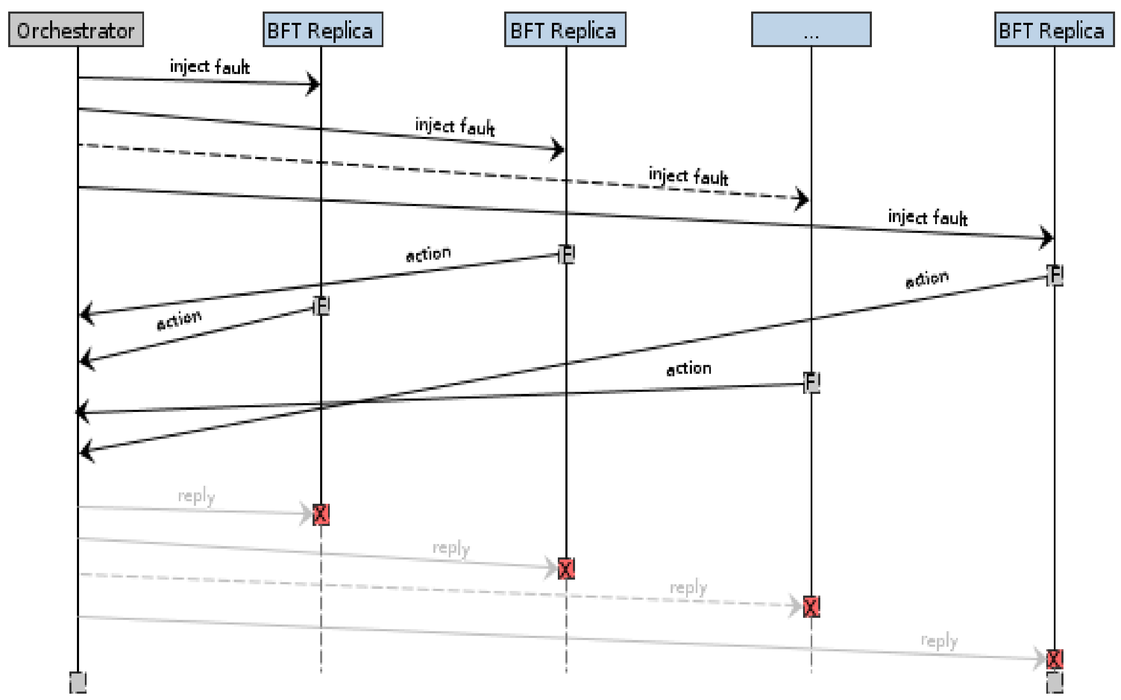

Case Study – BFT-SMaRt Protocol

Evaluation

- BFT-SMaRt - Byzantine Fault-Tolerance Protocol
  - Leader-based BFT supporting Dynamic Membership and Reconfiguration (FCUL)
  - 10 attacks designed and implemented, with special focus to the leader change sub-protocol

Final Result

Proposed fixes and tuned parameters allowed for a 10X recovery time improvement

Insights

- **Bugs**
  - Two low-level implementations bugs (underflow and overflow)
- **Proposed fix to handle timeout management**
  - Attacks on the leader could stall the protocol for more than 60s
- **Fixed parameters misconfiguration**
  - Protocol aborts and slow recovery caused by misconfigured default parameters. Provided tune-up values
- **Proposed fix to prevent suppression of low-level errors**
  - i.e., corrupted contents in a packet were silently discarded in the communication channel and no information was sent to the reconfiguration protocol

Intel Science & Technology Center for Cloud Computing

Carnegie Mellon University · Georgia Tech · intel · Princeton University · UC Berkeley · University of Washington