

An Overview of Existing Tools for Fault-Injection and Dependability Benchmarking in Grids

Sébastien Tixeuil¹ and William Hoarau¹ and Luis Silva²

¹*LRI - CNRS UMR 8623 & INRIA Grand Large,
Universit Paris Sud XI, France
Email: {tixeuil,hoarau}@lri.fr*

²*Departamento Engenharia Informtica, Univ. Coimbra,
Polo II, 3030-Coimbra, Portugal
Email: luis@dei.uc.pt*



CoreGRID Technical Report
Number TR-0041
October 5, 2006

Institute on System Architecture

CoreGRID - Network of Excellence
URL: <http://www.coregrid.net>

CoreGRID is a Network of Excellence funded by the European Commission under the Sixth Framework Programme

Project no. FP6-004265

An Overview of Existing Tools for Fault-Injection and Dependability Benchmarking in Grids

Sébastien Tixeuil¹ and William Hoarau¹ and Luis Silva²

¹LRI - CNRS UMR 8623 & INRIA Grand Large,
Universit Paris Sud XI, France
Email: {tixeuil,hoarau}@lri.fr

²Departamento Engenharia Informtica, Univ. Coimbra,
Polo II, 3030-Coimbra, Portugal
Email: luis@dei.uc.pt

CoreGRID TR-0041
October 5, 2006

Abstract

In this paper we review several existing tools for fault injection and dependability benchmarking in grids. We emphasize on the FAIL-FCI fault-injection software that has been developed in INRIA Grand Large, and a benchmark tool called QUAKE that has been developed in the University of Coimbra. We present the state-of-the-art and we explain the importance of these tools for dependability assessment of Grid-based applications and Grid middleware.

1 Introduction

One of the topics of paramount importance in the development of Grid middleware is the impact of faults since their probability of occurrence in a Grid infrastructure and in large-scale distributed system is actually very high. So it is mandatory that Grid middleware should be itself reliable and should provide a comprehensive support for fault-tolerance mechanisms, like failure-detection, checkpointing-recovery, replication, software rejuvenation, component-based reconfiguration, among others. One of the techniques to evaluate the effectiveness of those fault-tolerance mechanisms and the reliability level of the Grid middleware is to make use of some fault-injection tool and robustness tester to conduct some experimental assessment of the dependability metrics of the target system. In this paper, we will present and review several software fault-injection tools and workload generators for Grid Services that can be used for dependability benchmarking in Grid Computing. The ultimate goal of our common work is to provide some contributions for the definition of a dependability-benchmark for Grid computing and to provide a set of tools and techniques that can be used by the developers of Grid middleware and Grid-based applications to conduct some dependability benchmarking of their systems.

Dependability benchmarking must provide a uniform, repeatable and cost-effective way of performing experiments in different systems, mechanisms or components. Those three metrics can be achieved with the development of software tools that will be used in the process and in the definition of the dependability benchmark. According to [1] a dependability benchmark should provide the following components:

Workload: which represents the work the system must do during the execution of the benchmark;

Faultload: represents a set of faults and stressful conditions to emulate real faults that experienced in the real systems;

This research work is carried out under the FP6 Network of Excellence CoreGRID funded by the European Commission (Contract IST-2002-004265).

Measures: characterize the performance and dependability of the system under benchmark in the presence of the faultload when executing the workload;

Experimental setup and benchmark procedure: describes the setup required to run the benchmark and the set of rules that should be followed during the benchmark execution.

Note that the faultload that is generated here only considers benign faults that simply eventually exhaust resources during the execution of the benchmark. For actual fault tolerance testing, a more complex set of tests need to be carried out. Due to the inherent non-determinism of distributed systems (which include Grids), automatically testing against faults faces the problem of reproducibility of fault scenarios to check the resilience of the system. There are several components that characterize software fault-injectors in the specific case of grids:

Fault scenario: this represents the way faults are expected to occur in a particular run of the system;

Interposition system: this represents how faults are actually injected in the application being run (typically an interposition system between the operating system and the application under test);

Measures: this represents the performance and handling of faults when a particular fault scenario is executed.

In this paper we review several available (academic) tools for fault injection (Section 2) and dependability benchmarking (Section 3) for Grids. We present the two tools that have been developed by the two partners of WP4 from CoreGrid (INRIA and Univ. Coimbra) and we explain some of our ongoing projects with these tools.

2 Fault-injection Tools

When considering solutions for software fault injection in distributed systems, there are several important parameters to consider. The main criterion is the usability of the fault injection platform. If it is more difficult to write fault scenarios than to actually write the tested applications, those fault scenarios are likely to be dropped from the set of performed tests. The issues in testing component-based distributed systems have already been described and methodology for testing components and systems has already been proposed [2, 3]. However, testing for fault tolerance remains a challenging issue. Indeed, in available systems, the fault-recovery code is rarely executed in the testbed as faults rarely get triggered. As the ability of a system to perform well in the presence of faults depends on the correctness of the fault-recovery code, it is mandatory to actually test this code. Testing based on fault-injection can be used to test for fault-tolerance by injecting faults into a system under test and observing its behavior. The most obvious point is that simple tests (e.g. every few minutes or so, a randomly chosen machine crashes) should be simple to write and deploy. On the other hand, it should be possible to inject faults for very specific cases (e.g. in a particular global state of the application), even if it requires a better understanding of the tested application. Also, decoupling the fault injection platform from the tested application is a desirable property, as different groups can concentrate on different aspects of fault-tolerance.

Decoupling requires that no source code modification of the tested application should be necessary to inject faults. Also, having experts in fault-tolerance test particular scenarios for application they have no knowledge of favors describing fault scenarios using a high-level language, that abstract practical issues such that communications and scheduling. Finally, to properly evaluate a distributed application in the context of faults, the impact of the fault injection platform should be kept low, even if the number of machines is high. Of course, the impact is doomed to increase with the complexity of the fault scenario, e.g. when every action of every processor is likely to trigger a fault action, injecting those faults will induce an overhead that is certainly not negligible.

2.1 CECIUM

CECIUM [16] is a testing environment that simulates a distributed application and fault injections. The distributed execution is thus simulated on a unique machine with a unique addressing space. The source code of the tested application is not required.

The execution being simulated, the reproducibility of tests becomes trivial. However, in a real distributed system, very complex faults due to very complex scenarios may appear, and the model used by CECIUM may not include all possibilities due to the inherent limitations of simulations vs. actual software and hardware.

Another simulation based tool is MEFISTO [17] (Multi-level Error/Fault Injection Simulation Tool), a tool that permits to inject faults in VHDL model simulations.

2.2 DOCTOR

A few tools do exist to emulate the occurrence of faults in distributed applications. One of those tools is DOCTOR [4] (integrateD sOftware fault inJeCTiOn enviRonment), that permits to inject faults in real time systems. It is capable to inject faults and to synthesize the workload. It supports three kinds of faults: processor faults, memory faults and communication faults. The injected faults can be permanent, transient, or intermittent. DOCTOR has been implemented in the HEART real-time system. During the experiments, DOCTOR collects performance and reliability information, and a GUI is provided.

However, the kind of faults that can be generated is not sufficient for large scale distributed systems such as Grids. Indeed the main kind of faults is crash in those systems, and crashes are not available in DOCTOR. Moreover, the fault scenarios that can be designed can only be probabilistic (and using the GUI). While this is sufficient for quick quantitative tests, deterministic and reproducible fault injection capabilities are required for more complex fault scenarios.

2.3 ORCHESTRA

ORCHESTRA [5] is a fault injection tool. It allows the user to test the reliability and the liveliness of distributed protocols. A fault injection layer is inserted between the tested protocol layer and the lower layers, and allows to filter and manipulate messages exchanged between the protocol participants. Messages can be delayed, lost, reordered, duplicated, modified and new messages can be spontaneously introduced into the tested system to bring it into a particular global state.

The reception script and the sending script are written in Tcl language and determine which operations are to be performed on received/sent messages. These scripts are specified with state machines. Transitions in these state machines are driven by the type of the message, its contents, the history of received messages or other information that was previously collected during the test (e.g. local time, number of received messages, etc.). The message modifications are specified using a user-defined script. The resulting message is passed to the next layer of the protocol stack.

ORCHESTRA is a "Message-level fault injector" because a fault injection layer is inserted between two layers in the protocol stack. This kind of fault injector allows injecting faults without requiring the modification of the protocol source code. However, the user has to implement his faults injection layer for the protocol he uses. The expressiveness of the faults scenario is limited because there is no communication between the various state machines executed on every node. Then, as the faults injection is based on exchanged messages, the knowledge of the type and the size of these messages is required. In summary, ORCHESTRA is well adapted for the study of network protocols, but is too complex to use and not expressive enough to study distributed applications.

2.4 NFTAPE

The NFTAPE project [6] arose from the double observation that no tool is sufficient to inject all fault models and that it is difficult to port a particular tool to different systems. NFTAPE [6] provides mechanisms for fault-injection, triggering injections, producing workloads, detecting errors, and logging results. Unlike other tools, NFTAPE separates these components so that the user can create his own fault injectors and injection triggers using the provided interfaces. NFTAPE introduces the notion of Lightweight Fault Injector (LWFI). LWFIs are simpler than traditional fault injectors, because they don't need to integrate triggers, logging mechanisms, and communication support. This way, NFTAPE can inject faults using any fault injection method and any fault model. Interfaces for the other components are also defined to facilitate portability to new systems.

In NFTAPE, the execution of a test scenario is centralized. A particular computer, called the control host, takes all control decisions. This computer is generally separated from the set of computers that execute the test. It executes a script written in Jython (Jython is a subset of the Python language) which defines the faults scenario. All participating computers are attached to a process manager which in turn communicates with the control host. The control host sends commands to process managers according to the fault scenario. When receiving a command, the process manager executes it. At the end of the execution or if a crash occurs, the process manager notifies the control host by sending a message. All decisions are taken by the controller, which implies that every fault triggered at every node induces a communication with the controller. Then, according to the defined scenario, the controller sends a fault injection message to the appropriate process manager which can then inject the fault.

Although NFTAPE is modular and very portable, the choice of a completely centralized decision makes it very intrusive (its execution strongly perturbs the system being tested) if a considerable amount of resources is not dedicated to the controller. Scalability is then an issue if the controller is to manage a large number of process managers, and NFTAPE is run on regular clusters or grids (where there is no over-powered node that can take care of controlling the fault injector). Furthermore, implementing tests can end up in being very difficult for the user (the user may have to implement every component).

2.5 LOKI

LOKI [7] is a fault injector dedicated to distributed systems. It is based on a partial view of the global state of the distributed system. The faults are injected based on a global state of the system. An analysis is executed at the end of the test to infer a global schedule from the various partial views and then verify if faults were correctly injected (i.e. according to the planned scenario). Normally, injecting faults based on a global state of a distributed application leads to a high impact on its execution time. However, the technique used by LOKI is used to verify the validity of the injected faults while limiting their impact on the execution time.

In LOKI, every process of the distributed system is attached to a LOKI runtime to form a node. The LOKI runtime includes the code which manages the maintenance of the partial view of the global state, injects faults when the system arrives in a particular state, collects information about state changes and faults injections. LOKI proposes three execution modes: centralized, partially distributed and fully distributed. The definition of a scenario in LOKI is made by specifying the state machine used to maintain the local state, faults that are to be injected, and implementing a probe used to instrument the application in order to detect events (events are used to trigger faults). The user has to define fault identifiers and associate them to global states of the tested application, and then has to implement the probe by modifying the application source code. Calls to functions of the LOKI library must be inserted into the source code of the tested application to notify the LOKI runtime about events so that the appropriate state is reached. Also, to inject faults, the user has to implement the `injectFault()` function and insert it into the source code of the tested application. LOKI is the first fault injector for distributed systems that allows to inject faults based on a global state of the system and to verify if these faults were correctly injected. However, it requires the modification of the source code of the tested application. Furthermore, faults scenario are only based on the global state of the system and it is difficult (if not impossible) to specify more complex faults scenario (for example injecting "cascading" faults). Also, LOKI doesn't provide any support for randomized fault injection.

2.6 Mendosus

In [8] is presented Mendosus, a fault-injection tool for system-area networks that is based on the emulation of clusters of computers and different network configurations. This tool made some first steps in the fault-injection and assessment of faults in large distributed systems.

2.7 OGSA

In [9] is presented a fault-injection tool that was specially developed to assess the dependability of Grid (OGSA) middleware. This is the work more related with ours and we welcome the first contributions done by those authors in the area of grid middleware dependability. However, the tool described in that paper is very limited since it only allows the injection of faults in the XML messages in the OGSA middleware, which seems to be a bit far from the real faults experienced in real systems.

2.8 FAIL-FCI

FAIL-FCI [15] is a recently developed tool from INRIA. First, FAIL (for FAult Injection Language) is a language that permits to easily described fault scenarios. Second, FCI (for FAIL Cluster Implementation) is a distributed fault injection platform whose input language for describing fault scenarios is FAIL. Both components are developed as part of the Grid eXplorer project [10] which aims at emulating large-scale networks on smaller clusters or grids.

The FAIL language allows defining fault scenarios. A scenario describes, using a high-level abstract language, state machines which model fault occurrences. The FAIL language also describes the association between these state

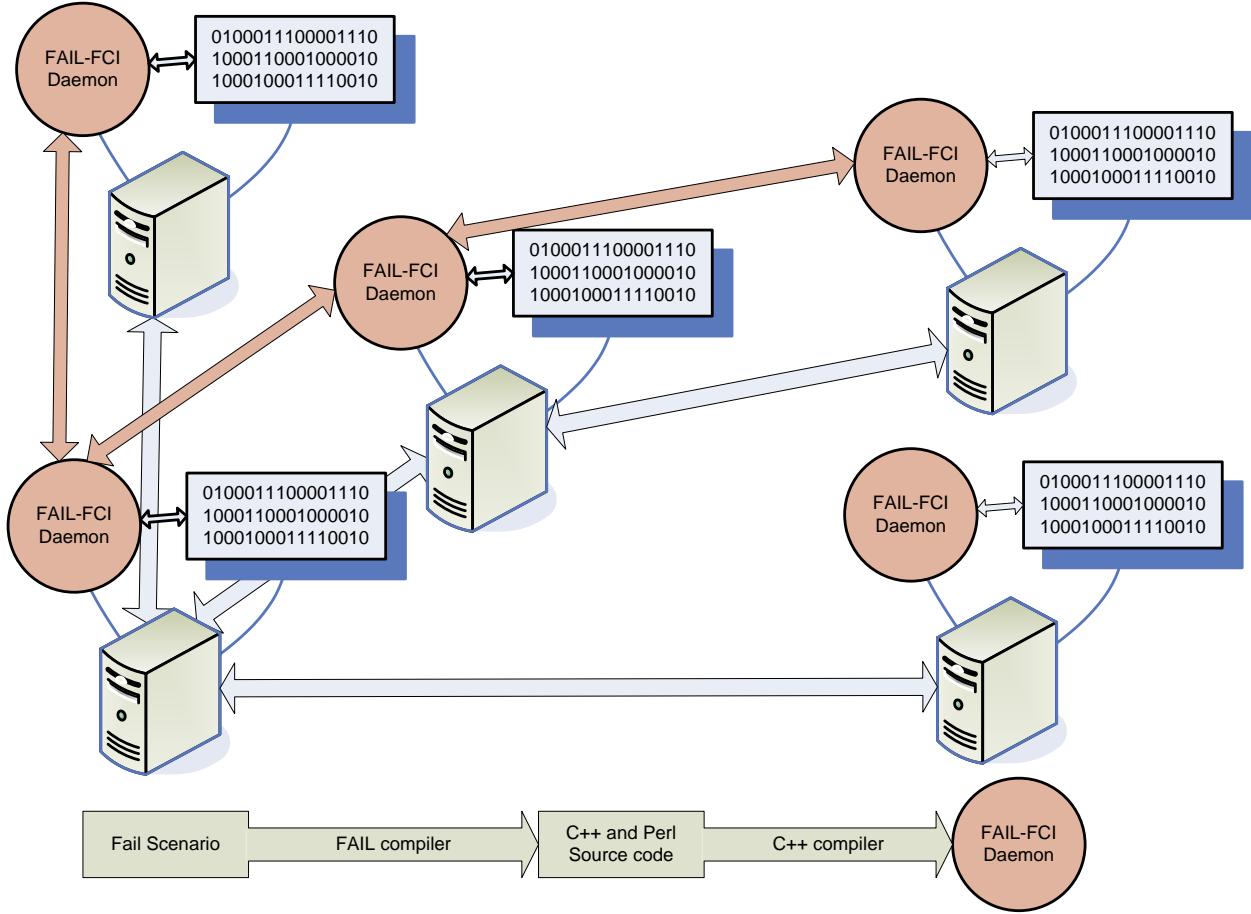


Figure 1: The FCI Platform

machines and a computer (or a group of computers) in the network. The FCI platform (see Figure 1) is composed of several building blocks:

- 1. The FCI compiler:** The fault scenarios written in FAIL are pre-compiled by the FCI compiler which generates C++ source files and default configuration files.
- 2. The FCI library:** The files generated by the FCI compiler are bundled with the FCI library into several archives, and then distributed across the network to the target machines according to the user-defined configuration files. Both the FCI compiler generated files and the FCI library files are provided as source code archives, to enable support for heterogeneous clusters.
- 3. The FCI daemon:** The source files that have been distributed to the target machines are then extracted and compiled to generate specific executable files for every computer in the system. Those executables are referred to as the FCI daemons. When the experiment begins, the distributed application to be tested is executed through the FCI daemon installed on every computer, to allow its instrumentation and its handling according to the fault scenario.

The approach is based on the use of a software debugger. Like the Mantis parallel debugger [11], FCI communicates to and from gdb (the Free Software Foundation's portable sequential debugging environment) through Unix pipes. But contrary to Mantis approach, communications with the debugger must be kept to a minimum to guarantee low overhead of the fault injection platform (in our approach, the debugger is only used to trigger and inject software faults). The tested application can be interrupted when it calls a particular function or upon executing a particular line of its source code. Its execution can be resumed depending on the considered fault scenario.

With FCI, every physical machine is associated to a fault injection daemon. The fault scenario is described in a high-level language and compiled to obtain a C++ code which will be distributed on the machines participating to the experiment. This C++ code is compiled on every machine to generate the fault injection daemon. Once this preliminary task has been performed, the experience is then ready to be launched. The daemon associated to a particular computer consists in:

1. a state machine implementing the fault scenario,
2. a module for communicating with the other daemons (e.g. to inject faults based on a global state of the system),
3. a module for time-management (e.g. to allow time-based fault injection),
4. a module to instrument the tested application (by driving the debugger), and
5. a module for managing events (to trigger faults).

FCI is thus a Debugger-based Fault Injector because the injection of faults and the instrumentation of the tested application is made using a debugger. This makes it possible not to have to modify the source code of the tested application, while enabling the possibility of injecting arbitrary faults (modification of the program counter or the local variables to simulate a buffer overflow attack, etc.). From the user point of view, it is sufficient to specify a fault scenario written in FAIL to define an experiment. The source code of the fault injection daemons is automatically generated. These daemons communicate between them explicitly according to the user-defined scenario. This allows the injection of faults based either on a global state of the system or on more complex mechanisms involving several machines (e.g. a cascading fault injection). In addition, the fully distributed architecture of the FCI daemons makes it scalable, which is necessary in the context of emulating large-scale distributed systems. FCI daemons have two operating modes: a random mode and a deterministic mode. These two modes allow fault injection based on a probabilistic fault scenario (for the first case) or based on a deterministic and reproducible fault scenario (for the second case). Using a debugger to trigger faults also permits to limit the intrusion of the fault injector during the experiment. Indeed, the debugger places breakpoints which correspond to the user-defined fault scenario and then runs the tested application. As long as no breakpoint is reached, the application runs normally and the debugger remains inactive.

2.9 Summary

The essential differences between the main aforementioned fault-injectors for grids appear in the following table.

Table 1: Comparison of Fault-injection Tools

| <i>Criteria</i> | <i>ORCHESTRA [5]</i> | <i>NFTAPE [6]</i> | <i>LOKI [7]</i> | <i>FAIL-FCI</i> |
|------------------------------|----------------------|-------------------|-----------------|-----------------|
| High Expressiveness | no | yes | no | yes |
| High-level Language | no | no | no | yes |
| No Source Code Modification | yes | no | no | yes |
| Scalability | yes | no | yes | yes |
| Probabilistic Scenario | yes | yes | no | yes |
| Global-state-based Injection | no | yes | yes | yes |

3 Dependability Benchmarking

The idea of dependability benchmarking is now a hot-topic of research and there are already several publications in the literature. The components of a dependability benchmark have been defined in [1].

3.1 State-of-the-art

In [18] is proposed a dependability benchmark for transactional systems (DBench-OLTP). Another dependability benchmark for transactional systems is proposed in [19]. Both benchmarks adopted the workload from the TPC-C

performance bench-mark. While [18] used software-based faults, the work described on [19] considered a fault-load based on hardware faults. The Special Interest Group on Dependability Benchmarking (SIGDeB), created by the IFIP WG 10.4 in 1999, released a preliminary proposal with a set of standardized classes for the classification of dependability in transactional database systems [20]. The goal is to help out the comparison of computer systems concerning four different dimensions: availability, data integrity, disaster recovery and security.

A dependability benchmark for operating systems was proposed by [21]. That benchmark was targeted for the study of the operating system robustness in the scenario of faulty applications. Another study about the behavior of the operating system in the presence of software faults in OS components was presented in [22].

The research presented in [23] addresses the impact of human errors in system dependability. In [24] is presented a methodology to evaluate human-assisted failure-recovery tools and processes in server systems. That methodology was illustrated with a case study of undo/redo recovery tool in email services. Another work was presented in [25] that focus on the availability benchmarking of computer systems. The authors propose a methodology, including single and multi-fault workloads, and they applied that methodology to measure the availability of software RAID systems in different operating systems.

Research work at Sun Microsystems defined a high-level framework that is targeted to availability benchmarking [26]. That framework decomposes availability in three main components: fault-maintenance rate, robustness and recovery. Within the scope of that framework, they have developed two benchmarks: one that addresses specific aspects of a system's robustness on handling maintenance events such as the replacement of a failed hardware component or the installation of software patch [27]; and another benchmark that is related to system recovery [28].

At IBM, the Autonomic Computing initiative is also developing benchmarks to quantify the autonomic capability of a system [29]. In that paper they have discussed the requirements of benchmarks to assess the self-* properties of a system and they proposed a set of metrics for evaluation. In [30] is presented a further discussion about benchmarking the autonomic capabilities of a system. The authors present the main challenges and pitfalls. In [31] is presented an interesting approach to conduct bench-marking of the configuration complexity. This is a valuable contribution since one of the main problems of current IT systems is the complexity of deployment and management. A benchmark for assessing the self-healing capabilities of a system was presented in [32]. Two metrics were introduced: (a) the effectiveness of the system to heal itself in the occurrence of some perturbations; (b) a measure of how autonomic that healing action was achieved. This paper has clear connections with the work we are conducting in the study of self-healing techniques for SOAP-based servers.

In [33] the authors present a dependability benchmark for Web-servers (Web-DB). This tool used the experimental setup, the workload and the performance measures specified in the SPECWeb99 performance benchmark. Web-DB defined measures in the baseline performance, in the performance in presence of faults and some dependability measures, like autonomy, availability and accuracy.

The dependability benchmark tool that is presented in this paper is targeted to Grid and Web-services. We are mainly interested in the study of potential software aging problems and the effectiveness of self-healing techniques like software rejuvenation.

3.2 QUAKE: A Dependability Benchmark Tool for Grid Services

QUAKE is a dependability benchmark tool for Grid and Web-services. The following subsections present the QUAKE tool and the relevant metrics for dependability benchmark.

3.2.1 Experimental Setup and Benchmark Procedure

The QUAKE tool is composed by the following components presented in Figure 2. The main components are the Benchmark Management System (BMS and the System Under Test (SUT). The SUT consists of a SOAP server running some Web/Grid service. From the point of view of the benchmark the SUT corresponds to a Web-based application server, a SOAP router and a Web-service. That Web-service will execute under some workload, and optionally will be affected by some fault-load. There are several client machines that invoke requests in the server using SOAP-XML requests. All the machines in the infrastructure are clock-synchronized using NTP. The application under test is not limited to a SOAP-based application: in fact, the benchmark infrastructure can also be used with other examples of client-server applications that use other different middleware technologies.

The Benchmark Management System (BMS) is a collection of software tools that allows the automatic execution of the benchmark. It includes a module for the definition of the benchmark, a set of procedures and rules, definition

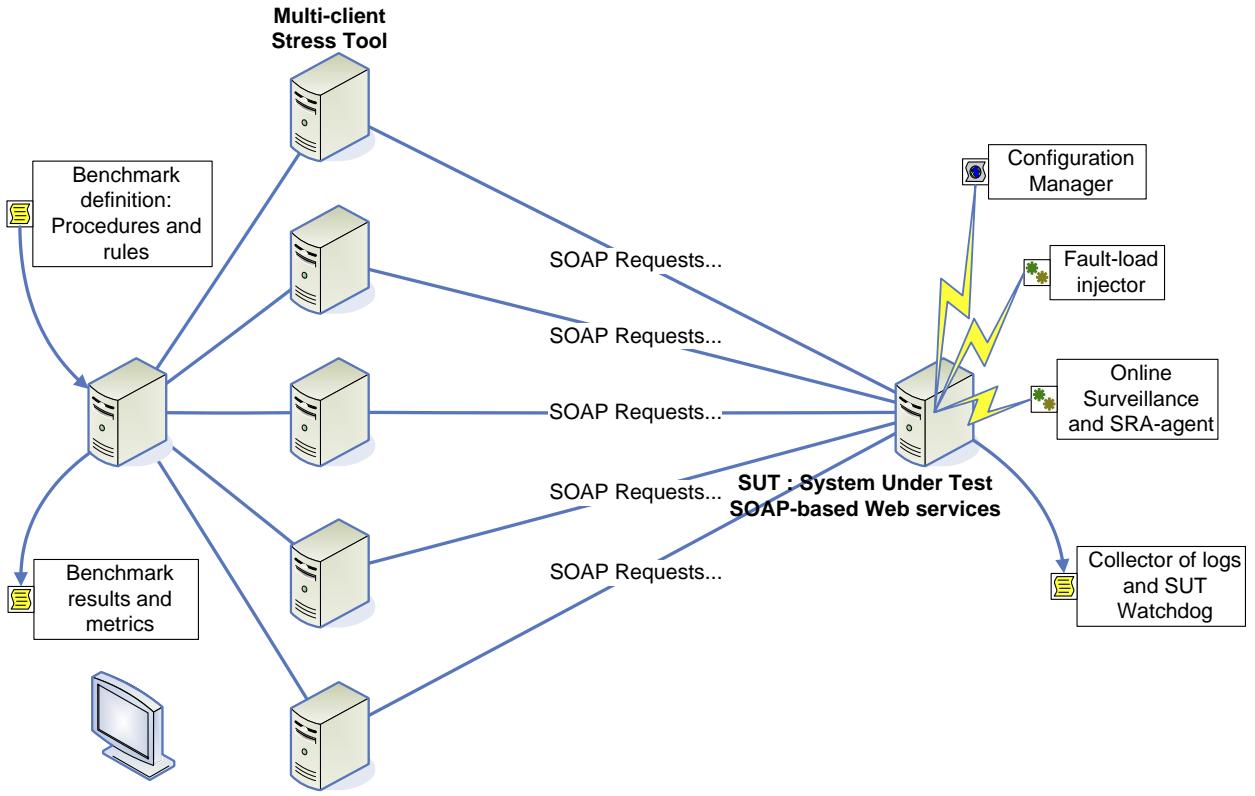


Figure 2: Experimental setup of the QUAKE tool

of the workload that will be produced in the SUT, a module that collects all the benchmark results and produces some results that are expressed as a set of dependability metrics. The BMS system may activate a set of clients (running in separate machines) that inject the defined workload in the SUT by making SOAP requests to the Grid Service. The execution of the client machines is timely synchronized and all the partial results collected by each individual client are merged into a global set of results that generated the final assessment of the dependability metrics. The BMS system includes a reporting tool that presents the final results in a readable and graphic format.

The Benchmark Management System (BMS) is a collection of software tools that allows the automatic execution of the benchmark. It includes a module for the definition of the benchmark, a set of procedures and rules, definition of the workload that will be produced in the SUT, a module that collects all the benchmark results and produces some results that are expressed as a set of dependability metrics. The BMS system may activate a set of clients (running in separate machines) that inject the defined workload in the SUT by making SOAP requests to the Web Service. The execution of the client machines is timely synchronized and all the partial results collected by each individual client are merged into a global set of results that generated the final assessment of the dependability metrics. The BMS system includes a reporting tool that presents the final results in a readable and graphic format.

The results generated by each benchmark run are expressed as throughput-over-time (requests-per-second in a time axis), the total turnaround time of the execution, the average latency, the functionality of the services, the occurrence of failures in the Web-service, the characterization of those failures (crash, hang, zombie-server), the correctness of the final results and the failure scenarios that are observed at the client machines (explicit SOAP error messages or time-outs).

From the side of the SUT system, there are four modules that also make part of the QUAKE benchmark tool: a fault-load injector, a configuration manager, a collector of logs with the benchmark results and a watchdog of the SUT system.

The configuration manager helps in the definition of the configuration parameters of the SUT middleware. It is absolutely that the configuration parameters may have a considerable impact in the robustness of the SUT system. By

changing those parameters in different runs of the benchmark it allow us to assess the impact of those parameters in the results expressed as dependability metrics.

The SUT system should also be installed with a module to collect raw data from the benchmark execution. This log data will be then sent to the BMS server that will merge and compare with the data collected from the client machines. The final module is a SUT-Watchdog that detects when a SUT system crashes or hangs when the benchmark is executing. When a crash or hang is detected the watchdog generates a restart of the SUT system and associated applications, thereby allowing an automatic execution of the benchmark runs without user intervention.

There is another external module that will be used but is related with the software rejuvenation study. This module includes two main components:

1. **An online Surveillance System:** this component is responsible for the online monitoring of some key parameters that may indicate the occurrence of software aging, namely: throughput (requests-per-second), average response time of requests and the memory usage at the server under test.
2. **A Software Rejuvenation Agent:** this module applies some predictive techniques based on time-series analysis to indicate in advance the probability of software aging in the server system. It is used to trigger rejuvenation actions to study the impact of the self-healing capability in the system-under-test.

In this particular study, this predictive system does not trigger automatically a rejuvenation action: it is mainly used to study the likelihood of software aging in some packages of SOAP middleware. The rejuvenation actions in this study are triggered by a threshold metric based on the SLA parameters of the Web-service. We are not only concerned in software aging that results in hang or crash situations. We are also considering the cases where the software aging manifests as a clear performance degradation that may violate the SLA for that service component, and surely result in profit-loss. In this case, we do apply a rejuvenation action.

3.2.2 Benchmark Procedure

The benchmark can run in three modes:

1. **Learning mode:** in this mode the main performance metrics of the Web-service application are collected. The idea is to take a "picture" of the Web-service parameters while it is "young" and to compare it in a long-running execution where it may get "old". In other words, we get the baseline performance parameters and define a SLA for that service with thresholds. During the long-run execution we collect the same parameters and when they deviate from the baseline (as defined in the SLA) more than a certain threshold the SRA agent will trigger a rejuvenation action.
2. **Workload-only:** in this mode the Web-service will be tested under stress situations of workload to see its behavior and the likelihood of software aging. No faultload is injected in this mode, so the Web-service will use all the available system resources in the SUT server. In this mode the Web-service parameters are monitored in order to detect deviations from the SLA defined from the baseline parameters.
3. **Workload and faultload:** in this mode, the behavior of the Web-service will be studied when we apply simultaneous a defined workload and some particular faultload in the system resources of the server machine.

3.2.3 Workload

Since we want to study the response of a Web-service in some stress situations we decided to include a list of different workloads that can be chosen at the beginning of each experiment. Currently, it is possible to use any of the following distributions:

- WL1** Continuous burst with a maximum workload;
- WL2** Steady-state distribution (N requests-per-second);
- WL3** Burst for X minutes, quiet for Y minutes;
- WL4** Surge load (fixed request rate and then a peak value in the load).

WL5 Ramp-up distribution: starts with an initial defined request-rate and increases at discrete intervals.

The client machines are all timely synchronized: they start executing at the same time and all of them use the same distribution for the request workload.

There are some other tools in the market that can be used for performance benchmarking of SOAP Web-services, namely: SOAtest [34] and TestMaker [35]. However those tools are mainly targeted for testing the functionality of the SOAP applications and to collect some performance figures. QUAKE has some fundamental differences: it is targeted to study the dependability attributes, it includes a different approach for the workload distributions, a fault-load module that will be explained next and is used to evaluate the self-healing capabilities of a SOAP server that provides support for some Grid or Web-services.

3.2.4 Faultload

The faultload injector does not inject faults directly in the software or in the SOAP messages like in the traditional fault injection tools. This injector is obviously oriented for the problem of software aging. We do not inject any software fault, any hardware fault or any operator fault. Instead, we just emulate resource exhaustion to accelerate the occurrence of software aging and to see the impact in the Web-service under test. Other types of stressful condition can also be used in future experiments. The fault-load is introduced by an external program that runs in the same server (SUT) and thus compete for the same system resources, consuming one or several of the following operating systems resources:

FL1 Consumption of memory using a ramp-up distribution;

FL2 Creation of a large number of threads;

FL3 Extensive usage of file-descriptors;

FL4 Consumption of database connections;

This list of system resources (the targets of our faultload) may be extended in future versions, but for the time being, those resources seemed to be the most critical ones in a Web-service application server. The main point here is that the faultload we "inject" is representative of real scenarios.

3.2.5 Measures

The resulting measures are computed from the data that is collected by the BMS system in the several test-runs. They are grouped in four categories:

1. Baseline performance measures;
2. Performance measures in the presence of stress workload distributions;
3. Performance measures in the presence of stress workload and injected fault-loads at the server side;
4. Dependability related measures.

The baseline performance measures correspond to the sustained response that is obtained when the SUT is "young". These figures are obtained in short-term runs to avoid the possible occurrence of software aging. The baseline measures that we have considered are the following:

THR: corresponds to the sustained throughput (requests-per-second);

ART: average response time for a request;

CR: this corresponds to the number of conforming requests that were performed with a response time lower than a maximum value (Max_RT). Any request to the Web-service that exceeds this value may generate a "time-out" and thereby it is important to measure the number of CR (conforming requests) in every test-run;

The performance measures in the presence of workload are THRw, ARTw and CRw, while the performance measures in the presence of workload and faultload are THRwf ARTwf and CRwf.

Independently from the performance measures, several system parameters are also collected during the benchmark runs to feed the knowledge base of the SRA agent, as explained before.

In addition to the performance-oriented measures, we also consider a set of dependability measures directly obtained by the QUAKE tool. Since we are particularly interested in studying software aging and the effectiveness of potential self-healing mechanisms we consider the following measures, taking into account the work published in [32] and [33]:

Integrity: reports the number of errors that were found in the data of the Web-service under test, in the presence of the workload and our faultload. At the end of each test-run we have a procedure for checking the integrity of the final state of the database. Since we are not injecting software/hardware faults it is expected that this metric will always report a value of 100

Availability: represents the time the system is available to execute the workload of each test run. The watchdog module that was included in the QUAKE tool is responsible for the assessment of this metric. In this metric we distinguish downtime when the server is down due to a failure from the downtime that may be introduced in a rejuvenation operation. In this latter case, we assume that in production cases there is a cluster with a load-balancer that sends the incoming requests to another server while a SOAP-server is being rejuvenated. It is worth noting that the availability measure is related to the experimental conditions of the QUAKE tool and does not represent field availability (unconditional availability). As it is typically the case for dependability benchmarking, this availability measured is meant to be used for comparison purposes.

Autonomy: this measure shows the need for manual intervention from the system manager to repair the SOAP server in the presence of a hang situation. Hang-type failures are much more expensive than clean crash-failures: while in the latter case the application server may recover autonomously with a restart operation, the hang-scenario requires the manual intervention of a system manager to kill processes and clean some system resources.

Self-Healing Effectiveness: this measure is only considered in some test-runs, where we want to study the effectiveness of the SRA to perform proactive rejuvenation of the SUT. More than the effectiveness of the technique, it measures the positive impact in the autonomy, availability and the performance of the Web-service.

3.2.6 Benchmark Properties

As was explained in [36] a benchmark should offer the following properties:

Repeatable: our QUAKE infrastructure obtains similar results when running several times with the same workload in the same SUT;

Portable: as will be presented later, our QUAKE tool allows the comparison of different applications in this domain, that maybe implemented with different SOAP packages and even with other communication middleware;

Realistic: the scenario portrait by the QUAKE tool represents typical Web-service applications and the harsh load conditions represented by QUAKE workload are actually common in this type of system. Furthermore, as the fault load does not inject any software or hardware faults (only consumes system resources), the QUAKE tool does not suffer from the representativeness difficulties that affect typical dependability benchmarks.

So, in summary the QUAKE tool is able to be used in other environments for testing traditional client-server and Grid-based applications where the subject of software aging and self-healing would be a point of concern. This section presented a general description of QUAKE.

4 Conclusions and Current Status

We reviewed several available tools for software fault injection and dependability benchmarking tools for grids. We emphasized on the FAIL-FCI fault injector developed by INRIA, and on the QUAKE dependability benchmark developed by the University of Coimbra.

The FAIL-FCI tool has so far only provided preliminary results on desktop grid middleware (XtremWeb) and P2P middleware (the FreePastry Distributed Hash Table). These results permitted to identify quantitative failure points in both tested middleware, as well as qualitative issues concerning the failure recovery of XtremWeb.

With the QUAKE tool we have been conducting the following experimental studies:

1. Evaluate the robustness of some existing SOAP servers (Apache Axis, JBoss, gSOAP, MS.NET);
2. Assess the reliability of different middleware for client/server applications;
3. Evaluate the effectiveness of some mechanisms for software rejuvenation;
4. Study the reliability of OGSA-DAI implementation;

After that last phase we are planning to assess the dependability level of further modules from the Globus middleware (GT4).

After the core of the tool set is properly packaged for clusters, our goal is to enable larger scale experiments by (i) using Grids instead of clusters, and (ii) integrate with emulation mechanisms.

Acknowledgements

This research work is carried out in part under the FP6 Network of Excellence CoreGRID funded by the European Commission (Contract IST-2002-004265).

References

- [1] P.Koopman, H.Madeira. "Dependability Benchmarking & Prediction: A Grand Challenge Technology Problem", Proc. 1st IEEE Int. Workshop on Real-Time Mission-Critical Systems: Grand Challenge Problems; Phoenix, Arizona, USA, Nov 1999
- [2] S Ghosh, AP Mathur, "Issues in Testing Distributed Component-Based Systems", 1st Int. ICSE Workshop on Testing Distributed Component-Based Systems, 1999
- [3] H. Madeira, M. Zenha Rela, F. Moreira, and J. G. Silva. "Rifle: A general purpose pin-level fault injector". In European Dependable Computing Conference, pages 199-216, 1994.
- [4] S. Han, K. Shin, and H. Rosenberg. "Doctor: An integrated software fault injection environment for distributed real-time systems", Proc. Computer Performance and Dependability Symposium, Erlangen, Germany, 1995.
- [5] S. Dawson, F. Jahanian, and T. Mitton. Orchestra: A fault injection environment for distributed systems. Proc. 26th International Symposium on Fault-Tolerant Computing (FTCS), pages 404-414, Sendai, Japan, June 1996.
- [6] D.T. Stott and al. Nftape: a framework for assessing dependability in distributed systems with lightweight fault injectors. In Proceedings of the IEEE International Computer Performance and Dependability Symposium, pages 91-100, March 2000.
- [7] R. Chandra, R. M. Lefever, M. Cukier, and W. H. Sanders. Loki: A state-driven fault injector for distributed systems. In In Proc. of the Int.Conf. on Dependable Systems and Networks, June 2000.
- [8] X. Li, R. Martin, K. Nagaraja, T. Nguyen, B.Zhang. "Mendosus: A SAN-based Fault-Injection Test-Bed for the Construction of Highly Network Services", Proc. 1st Workshop on Novel Use of System Area Networks (SAN-1), 2002
- [9] N. Looker, J.Xu. "Assessing the Dependability of OGSA Middleware by Fault-Injection", Proc. 22nd Int. Symposium on Reliable Distributed Systems, SRDS, 2003
- [10] <http://www.lri.fr/fci/GdX>

- [11] S. Lumetta and D. Culler. "The Mantis parallel debugger". In Proceedings of SPDT'96: SIGMETRICS Symposium on Parallel and Distributed Tools, pages 118-126, Philadelphia, Pennsylvania, May 1996.
- [12] G. Fedak, C. Germain, V. Nri, and F. Cappello. "XtremWeb: A generic global computing system". Proc. of IEEE Int. Symp. on Cluster Computing and the Grid, 2001.
- [13] I. Foster, C. Kesselman, J.M. Nick and S. Tuecke. "Grid Services for Distributed System Integration", IEEE Computer June 2002.
- [14] J. Kephart. "Research Challenges of Autonomic Computing", Proc. ICSE05, International Conference on Software Engineering, May 2005
- [15] William Hoarau, and Sbastien Tixeuil. "A language-driven tool for fault injection in distributed applications". In Proceedings of the IEEE/ACM Workshop GRID 2005, page to appear, Seattle, USA, November 2005.
- [16] G. Avarez and F. Cristian, "Centralized failure for distributed, fault-tolerant protocol testing," in Proceedings of the 17th IEEE International Conference on Distributed Computing Systems (ICDCS'97) May 1997.
- [17] E. Jenn, J. Arlat, M. Rimen, J. Ohlsson, and J. Karlsson, "Fault Injection into VHDL Models: The MEFISTO Tool", Proceedings of the 24th International Symposium on Fault-Tolerant Computing, June 15-17, 1994.
- [18] M. Vieira and H. Madeira, "A Dependability Benchmark for OLTP Application Environments", Proc. 29th Int. Conf. on Very Large Data Bases (VLDB-03), Berlin, Germany, 2003.
- [19] K. Buchacker and O. Tschaeche, "TPC Benchmark-c version 5.2 Dependability Benchmark Extensions", <http://www.faumachine.org/papers/tpcc-depend.pdf>, 2003
- [20] D. Wilson, B. Murphy and L. Spainhower. "Progress on Deining Standardized Classes of Computing the Dependability of Computer Systems", Proc. DSN 2002, Workshop on Dependability Benchmarking, Washington, D.C., USA, 2002.
- [21] A. Kalakech, K. Kanoun, Y. Crouzet and A. Arlat. "Benchmarking the Dependability of Windows NT, 2000 and XP", Proc. Int. Conf. on Dependable Systems and Networks (DSN 2004), Florence, Italy, 2004.
- [22] J. Dures, H. Madeira, "Characterization of Operating Systems Behaviour in the Presence of Faulty Drivers Through Software Fault Emulation", in Proc. 2002 Pacific Rim Int. Symposium Dependable Computing (PRDC-2002), pp. 201-209, Tsukuba, Japan, 2002.
- [23] A. Brown, L. Chung, and D. Patterson. "Including the Human Factor in Dependability Benchmarks", Proc. of the 2002 DSN Workshop on Dependability Benchmarking, Washington, D.C., June 2002.
- [24] A. Brown, L. Chung, W. Kakes, C. Ling, D. A. Patterson, "Dependability Benchmarking of Human-Assisted Recovery Processes", Dependable Computing and Communications, DSN 2004, Florence, Italy, June, 2004
- [25] A. Brown and D. Patterson, "Towards Availability Benchmarks: A Case Study of Software RAID Systems", Proc. of the 2000 USENIX Annual Technical Conference, San Diego, CA, June 2000
- [26] J. Zhu, J. Mauro, I. Pramanick. "R3 - A Framework for Availability Benchmarking", Proc. Int. Conf. on Dependable Systems and Networks (DSN 2003), USA, 2003.
- [27] J. Zhu, J. Mauro, and I. Pramanick, "Robustness Benchmarking for Hardware Maintenance Events", in Proc. Int. Conf. on Dependable Systems and Networks (DSN 2003), pp. 115-122, San Francisco, CA, USA, IEEE CS Press, 2003.
- [28] J. Mauro, J. Zhu, I. Pramanick. "The System Recovery Benchmark", in Proc. 2004 Pacific Rim Int. Symp. on Dependable Computing, Papeete, Polynesia, 2004.
- [29] S. Lightstone, J. Hellerstein, W. Tetzlaff, P. Janson, E. Lassetre, C. Norton, B. Rajaraman and L. Spainhower. "Towards Benchmarking Autonomic Computing Maturity", 1st IEEE Conf. on Industrial Automatics (INDIN-2003), Canada, August 2003.

- [30] A.Brown, J.Hellerstein, M.Hogstrom, T.Lau, S.Lightstone, P.Shum, M.P.Yost, "Benchmarking Autonomic Capabilities: Promises and Pitfalls", Proc. Int. Conf. on Autonomic Computing (ICAC'04), 2004
- [31] A. Brown and J. Hellerstein, "An Approach to Benchmarking Configuration Complexity", Proc. of the 11th ACM SIGOPS European Workshop, Leuven, Belgium, September 2004
- [32] A.Brown, C.Redlin. "Measuring the Effectiveness of Self-Healing Autonomic Systems", Proc. 2nd Int. Conf. on Autonomic Computing (ICAC'05), 2005
- [33] J. Dures, M. Vieira and H. Madeira. "Dependability Benchmarking of Web-servers", Proc. 23rd International Conference, SAFECOMP 2004, Potsdam, Germany, September 2004. Lecture Notes in Computer Science, Volume 3219/2004
- [34] Parasoft SOAtest: <http://www.parasoft.com>
- [35] PushToTest TestMaker: <http://www.pushtotest.com>
- [36] J.Gray (Ed). "The Benchmark Handbook", Morgan-Kaufmann Publishers, 1993