

Operational Concepts

The following sections describe concepts to operate the E2E Bridge in a reliable and scalable manner. How to do this depends on the use case. The two most common use cases are **online services** and **batch processes**. Both scenarios must take into account different requirements since online services may come under heavy load depending on the number of parallel running clients whereas batch processes are normally not run in parallel but frequently process big files.

High Availability

High availability is a main point for integration systems. The E2E Bridge addresses high availability on different layers:

Robustness of the Service

Service developers should address the topic of reaction on “expected” errors (e.g. non-validation inputs) or temporary application outages at design time already. Approaches to achieve this are:

- **Time-outs or retries**
Even on top level (BPMN) time-outs or retries can be defined with modelling xUML services. For more information, refer to [BPMN execution on the Bridge](#) (esp. [BPMN Error Handling](#)) and [Persistent States and Signals](#) (esp. [Persistent State Error Handling](#)).
- **Store & forward approach**
This means, the E2E Bridge is always ready to receive messages, even when the final recipient is temporarily unavailable. In Bridge context, this store & forward approach can be implemented by two means:
 - By using the concept of UML state machines, the E2E Bridge achieves persistence by storing any transaction or state transition in a database. For more information refer to [Persistent States and Signals](#).
To use this approach, you need a database installation (see [Database Management System](#)).
 - By using an external messaging / queuing system: the E2E Bridge supports JMS. For more information refer to [JMS Adapter](#).
To use this approach, you need a JMS installation.

Self-recovery of the E2E Bridge

The E2E Bridge has the built-in capability to recover itself in case of unexpected errors:

- The Bridge can re-connect automatically after temporary connection loss, e.g. database outage, SAP connection loss, message queue connection loss.
- In case of fatal service errors or upon system reboot, the E2E Bridge can re-start service instances automatically. See [Preferences of an xUML Service](#) for more information.

Load Balancing

Service clients may create heavy load for service providers depending on the number of parallel running clients. A common way of distributing this load is using more than one node to host the services. Then, a load balancer will distribute the load among the identical services. The E2E Bridge can be integrated into any existing load-balancing infrastructure.

- **Active/active scenario**
Two redundant services are served by a [load balancer](#). As a general guideline, one server should be able to handle the full load alone, so that the load does not endanger continuity in case of an outage of the redundant server. Redundant online services usually share one highly available database.
- **Active/passive scenario**
A failover server takes over processing, if the primary server has a downtime. For batch services the active service is usually monitored by the redundant server. Should the corresponding services on the primary server have a downtime, a failover is done. This failover can be done by an underlying Cluster Software, or by the application itself. Refer to [Batch Processing](#) for more information on how to switch processing to a second node.

Refer to [Load Balancing](#) for a detailed description of a load balancing scenario.

Database Management System

The E2E Bridge persists any business data in an external database. For each service, the developer can decide on whether he wants to use a central fully-fledged RDBMS system or to rely on a small footprint, file based SQLite DB. SQLite provides for very high performance and scalability, but – as it is file-based - does not allow parallel access by different processes or threads.

On this Page:

- [High Availability](#)
 - [Robustness of the Service](#)
 - [Self-recovery of the E2E Bridge](#)
 - [Load Balancing](#)
 - [Database Management System](#)
- [Disaster Recovery](#)
- [Additional Topics](#)
 - [Monitoring](#)
 - [Security](#)
 - [Configuration Management](#)

Related Pages:

- [Batch Processing](#)
- [Configuration Management](#)
- [Load Balancing](#)
- [Secure Bridge Setup](#)

In scenarios that include high availability and disaster recovery concerns, the way to go is to use a full RDBMS system. We fully support Oracle®, MySQL® and MS SQL® databases. With Oracle, we even support clustered database setup and cluster failover handling as described on [Persistent State Using Clustered Oracle Database](#).

For separation of concerns, a dedicated DB schema per process is recommended. Within this schema, the E2E Bridge will basically create tables and manage their contents.

For more information on supported databases and how to install them, refer to [Installing and Configuring Database Access](#).

Disaster Recovery

The E2E Bridge uses an external RDBMS to store any application-relevant dynamic data, e.g. data retrieved by external applications, data calculated or collected by the services, status information about current process instances that are being handled. Considering that the E2E Bridge supports implicit and explicit transaction handling, this way all important business data is persisted in the database at any time. Business data recovery and continuity therefore is a question of the database's disaster recovery routine, which the E2E Bridge neatly integrates into.

Besides the database, the E2E Bridge also uses the file system to persist data (e.g. log files, infrastructural settings). The replication of this data is usually not business critical and thus irrelevant for disaster recovery, however an integration into the customers file backup mechanisms is recommended.

In general, for disaster recovery, a mirrored counterpart of the active productive system is recommended. For detailed information regarding high availability [see above](#). Refer to [Bridge Backup](#) for more information on how to take a Backup of an E2E Bridge installation and all deployed services.

Additional Topics

Monitoring

The Bridge monitors all deployed services. If a service writes a log entry of type ERROR or FATAL, or if a service terminates unexpectedly (crashes), the Bridge can call a monitoring service with all information found in the log file.

Refer to [Monitoring Load Balanced Nodes](#) for more information on the monitoring concept of the E2E Bridge.

Security

Security requirements can apply to the E2E Bridge itself, or to services running on the E2E Bridge.

- Refer to [Secure Bridge Setup](#) and [BRIDGE Hardening](#) for information on how to setup your Bridge installation within your enterprise infrastructure in a secure way.
- Refer to [Security Model](#) for information on how the E2E Bridge and E2E Builder support security during service development.

Configuration Management

Concerning service development, we recommend to deploy service repositories in a development, test, and productive environment. This development scenario is explained on [Configuration Management](#).