



Best practices

Defining your own EGO service to add High Availability capability for your existing applications

Leo Lin

*IBM Systems & Technology Group,
Software Defined Systems
Advisory Software Developer-Architect*

Contents

Configuring a reclaim grace period for Platform EGO	3
Prerequisites.....	3
Preparation.....	3
Steps	3
Notices	6
Trademarks	7

Configuring a reclaim grace period for Platform EGO

As the base for IBM Platform Symphony, IBM Platform EGO provides High Availability (HA) functionality to Platform Symphony when Platform Symphony daemons run as a Platform EGO service. A Platform EGO service has following HA capabilities:

- It can be restarted automatically if it dies unexpectedly.
- It be automatically restarted on another hosts if the running host is down

Users are able to define their own Platform EGO services to add HA capability to their existing applications.

The following steps describe how to define a simple Platform EGO service.

Prerequisites

You have knowledge of EGO consumer and resource groups

Preparation

1. Install the binaries and libraries that is used by `dummy` in all hosts that might run it. The Platform EGO service currently doesn't support provision of an application on demand.
2. If there is persistent state data, save the data to a network file system so that it can be access by all hosts that can potentially run `dummy`.

Steps

These steps describe how to keep an existing application (named `dummy` in the example)running even the host is down.

1. Go to the `$EGO_CONFDIR/../../eservice/esc/conf/service` directory and make a copy of the service definition template file `service.xml.TMPL` to any name with `.xml` extension. The following steps use the name `newserv.xml`.
2. Modify `newserv.xml`:
 - `<sc:MaxInstances>`: How many instances (processes) this service can start. For `dummy`, it has only one instance. Change it to 1.
 - `<sc:ControlPolicy>`
 - `<sc:StartType>`: If you want the service start automatically, change it to `AUTOMATIC`. Otherwise, leave it as `MANUAL`. You need either start by `egosh` or through the Platform Management Console.
 - `<sc:Dependancy>`: This is for a case when this service has dependency on another service. It is not applied here. Remove this tag.
 - `<sc:AllocationSpecification>`:
 - `<ego:ConsumerID>`: the consumer for this service. You can leave it as it is if you just want to have a try.
 - `<sc:ActivityDescription>`:
 - `<ego:Command>`: define the command of the service.

Such as:

```
<ego:Command>/opt/dummy_service/dummy -d network_file_dir</ego:Command>
```

Note: You can define different commands for different OS types (defined by `<ego:Attribute for name="hostType">`). `HostType` defined as "all" is applied if no corresponding OS type matches. Assume `dummy` is run on 64 bit Linux, define the command for `hostType="all"`.

3. Restart the master host.

Now that the service definition is done. Restart the master host to let the definition take effect.

```
egosh ego restart $MASTER_HOST
```

4. After the service is started, power off the running host and observe how the service will fail over to another machine.

For advanced users

Locate the service:

Platform Symphony default installation includes a DNS service named **ServiceDirector**. By **configuring ServiceDirector** and starting it, you can query DNS by the service name for the host where it is running. It is particularly useful if the service accepts incoming connections so that the client can always get the IP of the service host even if it fails over to another host.

The process of **ServiceDirector** is the "named". Its configuration file can be found at `$EGO_CONFDIR/../../eservice/esd.conf/named/conf`. Platform EGO automatically registers all services under the `.ego` domain.

For example: query the address of a Platform EGO service named `RS` :

```
$ dig RS.ego
; <<>> DiG 9.8.2rc1-RedHat-9.8.2-0.17.rc1.el6 <<>> RS.ego
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 9789
;; flags: qr aa rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 1, ADDITIONAL: 1

;; QUESTION SECTION:
;RS.ego. IN A

;; ANSWER SECTION:
RS.ego. 0 IN A 172.26.0.1
;; AUTHORITY SECTION:
ego. 0 IN NS egonameserver.ego.

;; ADDITIONAL SECTION:
egonameserver.ego. 0 IN A 172.26.0.1
```

```
;; Query time: 0 msec
;; SERVER: 127.0.0.1#53(127.0.0.1)
;; WHEN: Fri Nov 22 19:09:52 2013
;; MSG SIZE rcvd: 84
```

Define a service with multiple service instances

Configure `<sc:MaxInstances>` with a value `> 1`, to start multiple instances of one service provided there are enough available slots.

- One instance per host: In some use cases, you want to start multiple instances, but need those instances to run on different hosts. For this case, define `<sc:MaxInstancePerHost>` to 1:
 - a) Open service definition file in a text editor. (newserv.xml)
 - b) Add a new line after the definition of `<sc:MaxInstances></sc:MaxInstances>`

```
<sc:MaxInstancesPerHost>1</sc:MaxInstancesPerHost>
```
 - c) Restart the master host.

With this configuration, only once instance is started on each host.

- Multiple instances per host: If `MaxInstances` is greater than 1 and `MaxInstancesPerHost` is not defined, there likely will be multiple instances in any host. For a server that needs per instance resources (such as TCP listening ports, work directories, etc.), the instance `seqno` is provided as an environment variable `EGOSC_INSTANCE_SEQNO` for the process of the instance. For example:

Service `dummy` has 5 instances. Each instance requires its own working directory that can be passed as command line parameter with option `-d`.

Write a wrapper script `dummy.sh`:

```
mkdir /opt/dummy_service/work/${EGOSC_INSTANCE_SEQNO}
/opt/dummy_service/dummy -d ${EGOSC_INSTANCE_SEQNO}
```

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

Without limiting the above disclaimers, IBM provides no representations or warranties regarding the accuracy, reliability or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The use of this information or the implementation of any recommendations or techniques herein is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Anyone attempting to adapt these techniques to their own environment does so at their own risk.

This document and the information contained herein may be used solely in connection with the IBM products discussed in this document.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE: © Copyright IBM Corporation 2013 All Rights Reserved.

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml

Windows is a trademark of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.