



RisAris

Application Usage
Governance:
The SOA Gateway -
Central Governance
of your run time systems



Summary

This is a follow on document to the document 'SOA Gateway Governance – Lifecycle' which describes how the SOA Gateway can govern the lifecycle of a service. Once the service is deployed and operational, it is necessary to govern its usage. This document explores what data needs to be recorded to control this and how the SOA Gateway enables organizations to do this.



About the Author

John Power is a driving force behind the development of the [SOA Gateway](#) and Managing Director of Risaris Limited. With over 25 years experience in the software industry, John has delivered complex integration projects in Software AG, Delta Airlines, Boston University and Morgan Stanley.

Contents

1	Introduction.....	4
2	Definitions	4
3	Governance Data.....	4
4	SOA Gateway and Usage Governance.....	6
5	Conclusions.....	10
6	About The SOA Gateway.....	11

V.1.122/01/09

1 Introduction

Service in the context of this document is as defined in its sister document 'SOA Gateway Service Governance – Lifecycle Management'. When a service has been created and is deployed, there are a number of things that may need to be understood about the usage of service.

1. For audit purposes, it may be necessary to know who has used a specific service and when.
2. It may be necessary to log the data provided to the service and the data returned by the service to record what the user did.
3. It may be necessary to understand who is using a service on specific dates or even within a time range.
4. When a service has been deprecated (i.e. replaced by a newer version of the service), it is necessary to understand who is still using the service so that they may be upgraded to the newer level of the service and the older version retired.
5. For capacity planning purposes, it is necessary to understand how often a service is being called.
6. In order to comply with Service Level Agreements (SLAs), the behaviour of the service under normal circumstances must be understood so that a commitment can be made to a SLA and it can be detected when a service is performing outside its normal band.

This document describes what data is available to Govern services and how the SOA Gateway can make this available to an organization.

2 Definitions

The following terms are used in this document:

- SOA Gateway server session: This represents a single execution of a SOA Gateway from the time the server is started until it is terminated.

3 Governance Data

This section describes the data that is available to enable an organization to properly govern their various services based on their core software assets.

3.1 Identity

The identity of the user of a service is one of the more important aspects of Governance. The identity of the user could be:

1. The system userid of the user accessing the system meaning the userid of up to 8 bytes that many users have to access online systems.
2. The userid from the SSL Certificate if used to identify the user.
3. In some cases, the identity may be of a specific system which provides services to users but at least it will be clear where the request is coming from. If there is a requirement to see who is actually issuing the request, such systems would have to be changed to pass up the correct identity.
4. There may be anonymous services where this is not available, however, in such cases, the service is unlikely to require to be monitored as closely. If it is, it should be reviewed.

3.2 What time a service was used

The local server date and time that the service was used. While the request may originate in a different time zone, the local system must work to a standard date and time zone.

3.3 Performance Metrics

In order to work to SLAs and determine if services are working outside of their agreed characteristics, the following data may be collected for such services over a specific server time frame:

1. The number of times the service has been successfully called.
2. The number of times the service has been unsuccessfully called.
3. The total response time for successful calls to the service.
4. The lowest ever response time for the service.
5. The highest ever response time for the service.

A timeframe in the above context could relate to how the organization works. It could be:

1. For a given period of time (e.g. every hour, every 24 hours, every week or every month) with statistics being written and reset at a specific time which coincides with some business cycle.
2. For the duration of the SOA Gateway server session

Note that it may also be useful for some services to log relevant metrics (e.g. response time) for each invocation of the service, however, this would lead to a large amount of data being generated and should only be necessary under specific circumstances.

3.4 What was requested of the service

What data was provided to the service by the caller. This can be necessary where access to very sensitive information must be logged based on who accessed the service and what they asked for.

3.5 What was returned by the service

What data was returned to the caller by the service. Again, this can be necessary for very sensitive services where access to data must be carefully monitored.

3.6 Location of service user

The location from which the request came from will normally take the form of an IP address. This will be of more relevance within an Intranet where IP addresses will generally be fixed and known than in the wider Internet as it is difficult to prove beyond a reasonable doubt where the IP address was allocated at a specific time.

4 SOA Gateway and Usage Governance

The SOA Gateway enables the collection of usage information for services and can make this available to various governance software stacks.

4.1 Scope of Collection

There are two levels of collection available within the SOA Gateway. The first relates to global information that may be required and applies to all services and versions of the service in a SOA Gateway instance. The second relates to individual services and how much governance information is collected about each.

4.1.1 Global Collection

This type of collection occurs for global events that the installation may wish to know about which include the following which may be turned on or off as the organization wishes:

- Invocation of deprecated services. One record will be written for each deprecated service that is used during a SOA Gateway server session even if it is invoked multiple times. If the installation wishes to determine who is invoking this service, the service level definitions may be modified to get more comprehensive information about the usage.
- Out of Service Level Agreement (SLA) conditions for a service. This will be written when a service responds outside of the pre assigned service level. These events will be triggered every 1 minute with the count of times it has occurred within this 1 minute timeframe for a specific service. This is to prevent floods of messages for a service that is experiencing a problem and thus each service request will be outside of the SLA.
- Service failure. This will be written when a service does not respond normally to a client. In other words, a SOAP fault or other error is raised based on the invocation of the service. These events will be triggered every 1 minute with the count of times it has occurred within this 1 minute timeframe for a specific service. This is to prevent floods of messages for a service that is experiencing a problem and thus each service request will be outside of the SLA.

- Security errors. Whenever a security error is raised while running a service, this will be logged as a global event.

4.1.2 Service collection

The SOA Gateway can be configured to collect data or not at the service name level. This represents the logical name of the service which may have different versions of the service active at any one time in a given SOA Gateway server instance. The following information may be collected about each service.

- No governance information about the service at all.
- Summary performance information about the service. This will make information about the performance characteristics about the service available periodically. Once the information has been made available, the statistics will be reset in advance of the next period.
- Detailed performance information about the service. This will record the performance data (e.g. response time) for each request made to the service. For a heavily used service, this has the potential to generate a lot of information.
- Detailed access information about the service. This will record the access data (e.g. date and time of access, identity of caller, location of caller etc.) for each request made to the service. For a heavily used service, this has the potential to generate a lot of information.
- Request information provided to the service by the caller. This will record the message data provided to the service for each request made to the service. For a heavily used service, this has the potential to generate a lot of information.
- Response information returned by the service to the caller. This will record the message data returned by the service to the caller for each request made to the service. For a heavily used service, this has the potential to generate a lot of information.

4.1.3 Controlling Collection

Collection will be controlled by higher level defaults for the entire SOA Gateway server instance. It may then be overridden at each service level to override the default.

4.2 Collection of data

Data will only be collected for a given service when appropriate options are explicitly set or defaulted. The data will be collected in control blocks closely associated with the service block to which they relate for performance reasons. The data will be collected in an internal format as follows:

- Summary data will be held in a control block associated with the service with which it is related. This data will then be written based on the following events:
 - o On request from an external Governance server requesting this data.

- Prior to the service definition being deleted from the running system due to lack of use.
- Prior to the service being deleted from the configuration.
- Prior to the service being marked deprecated due to a current version of the service being deployed.
- Based on the global statistics settings for the SOA Gateway service.
- Detailed information will be collected individually and queued for later processing. Note the following about these events:
 - There will be an internal limit on the number of events that can be on the queue at any one time. If this is exceeded, events will be lost until the queue has been cleared.
 - The queue will be processed by an internal daemon process which will deal with the detailed events based on the configuration parameters for the server.

4.2.1 Writing Data to the Governance Component

The SOA Gateway server daemon may be configured to push information to a Governance component in three different ways:

1. It may be written to a file local to the SOA Gateway server for later processing.
2. It may be written to a REST service in which case the URL for the REST Service must be provided in the configuration.
3. It may be written to a SOAP service in which case the endpoint for the SOAP Service must be provided in the configuration.

Once the information has been processed, it will be reset or deleted as appropriate in the SOA Gateway server. Note that the configuration may differentiate between how summary and detailed information is handled as there are different operational requirements for detailed information due to the resources required to process it.

4.2.2 Pull from buffer or local file

The data may also be pulled from the SOA Gateway server as follows:

- Summary information may be pulled from the SOA Gateway server at any time.
- Detailed information may also be pulled from the SOA Gateway server, however, if this is not done often enough, resources may be unavailable and some events will be lost. In this case, it may be necessary to write the detailed information to a file local to the SOA Gateway server for download on request from the Governance technology.

4.3 Processing of Data

The Governance data produced by the SOA Gateway will all be in XML format for which a schema definition (i.e. an XSD) will be available. This will describe the data and format of the data being provided by the SOA Gateway and the Governance solution being used by an installation must have the ability to be configured to process the various types of

Governance data records from the SOA Gateway in some sort of intelligent manner. This data may then be used in the following ways:

1. To create dash boards based on the performance of each SOA Gateway and enable the installation to drill down to the performance of each individual service in that instance.
2. To monitor SLAs to ensure that services are operating within their agreed SLA.
3. To log access to a particular service or set of services.
4. To monitor usage of services that have been deprecated to determine where they are still in use.
5. To log what type of request data is being sent to specific users.
6. To log what type of response data is being sent to specific users.

With all of the above, due to the data being in XML format, it will be possible to create triggers based on the content of any of the XML nodes in the messages being processed. In addition, if logs need to be processed after the event, there are many tools available which can help with the processing of XML documents.

Finally, as all of the data from multiple SOA Gateway nodes will have the same format, the data can be merged and processed as a single unit if appropriate.

5 Conclusions

The SOA Gateway provides an ideal environment to govern the usage of your services exposing your core assets. This ensures that:

- All services are measured and logged in a consistent way based on policies set by the installation.
- Consistent recording of data means that realistic comparisons of the characteristics of services are possible
- The governance data provided by the SOA Gateway may be provided to any tool. It is up to an organization to choose what tool they wish to use as there are a number of excellent products on the market that do this very well and can tie SOA Gateway data in with data from other components of a transaction.
- It provides a perfect complement to the lifecycle governance provided by the SOA Gateway.

6 About The SOA Gateway

The SOA Gateway is a cost effective software tool to:

Access data faster...

It enables access to data from a wide range of database languages (ADABAS, MySQL, DB2, VSAM, Oracle etc.) without server side code, or expensive middleware.

Access business logic easier...

The SOA Gateway enables easy access and re-use of valuable business logic available in CICS, COBOL, C, NATURAL and many other languages and environments.

The [SOA Gateway](#) is developed by [integration specialists Risaris](#) Limited.

For a free trail of the [SOA Gateway](#), please visit:

http://www.soagateway.com/html/registration_form.php

This document is distributed for information purposes only and does not form part of or constitute an agreement with Risaris Ltd. Although Risaris Ltd. uses reasonable efforts to include accurate and up-to-date information in this document, Risaris makes no warranties or representations as to its accuracy. Risaris Ltd. may also make improvements and/or changes to this document at any time without notice. The various approaches outlined in this document are put forward in good faith, but it remains possible that individual results may vary. For that reason and in accordance with standard practice, readers are encouraged to test any materials developed on the basis of this paper before putting them into productive use.

More whitepapers at:

http://www.soagateway.com/html/industry_papers.html

See Terms & conditions at www.soagateway.com © Risaris Limited 2009.