

Statically and Dynamically verifiable SLA metrics

Violet Ka I Pun

Department of Informatics
University of Oslo

(Elena Giachino¹, Stijn de Gouw², Cosimo Laneve¹ and Behrooz Nobakht²)

¹ University of Bologna

² SDL Fredhopper

COST ARVI meeting





<http://www.envisage-project.eu>

Motivation




Motivation






Service Level Agreements (SLAs)

- ▶ Legal contracts between service providers and customers
- ▶ Define the quality of service
- ▶ **Informal**



Service Level Agreements (SLAs)

- ▶ Legal contracts between service providers and customers
- ▶ Define the quality of service
- ▶ **Informal**



Service Level Agreements (SLAs)

- ▶ Legal contracts between service providers and customers
- ▶ Define the quality of service
- ▶ **Informal**

Motivation – The gap between SLAs and services



- ▶ Difficult to assess whether a service complies with an SLA or not
- ▶ **Waste of resources and money**

Motivation – Bridge the gap between SLAs and services

Metrics:

- ① A means to specify QoS in SLAs
 - Define the boundaries and margins of **errors** wrt. cloud services
- ② Usage:
 - **Static time**: code evaluation
 - **Runtime**: service monitoring, balancing, or remediation

Motivation – Bridge the gap between SLAs and services

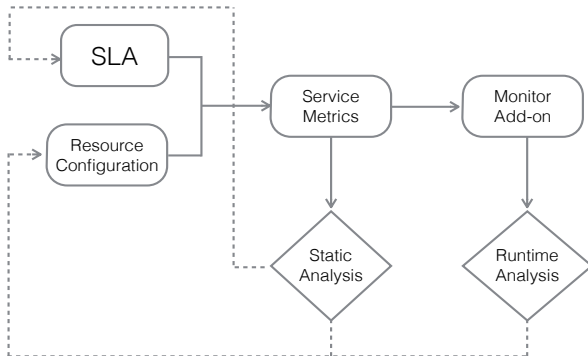
Metrics:

- ① A means to specify QoS in SLAs
 - Define the boundaries and margins of **errors** wrt. cloud services
- ② Usage:
 - **Static time**: code evaluation
 - **Runtime**: service monitoring, balancing, or remediation

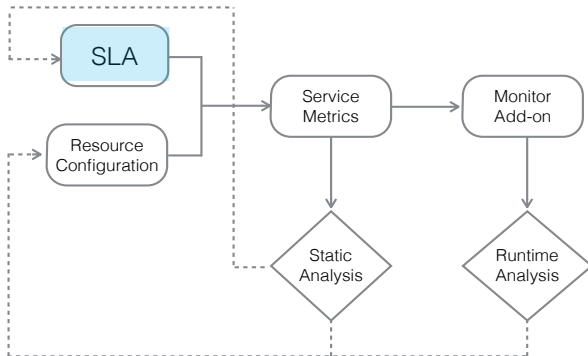
Our approach:

- ① Metric functions:
 - A simple **formal** description of SLA
- ② A mathematical framework
 - **Statically** derives the SLA quality levels from programs
 - **Dynamically** monitors service behaviours

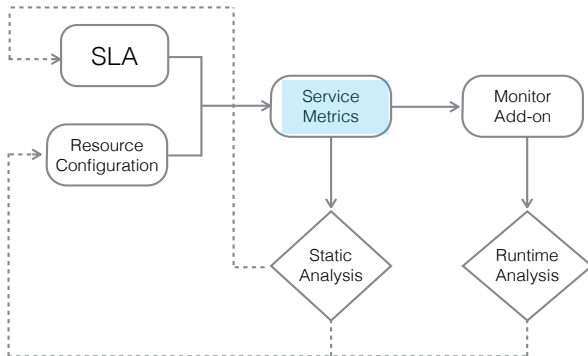
The Analysis Flow



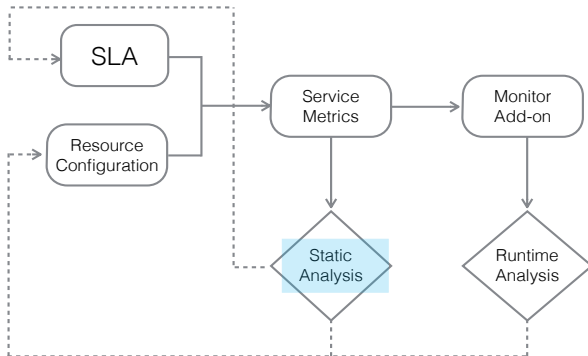
The Analysis Flow




The Analysis Flow




The Analysis Flow



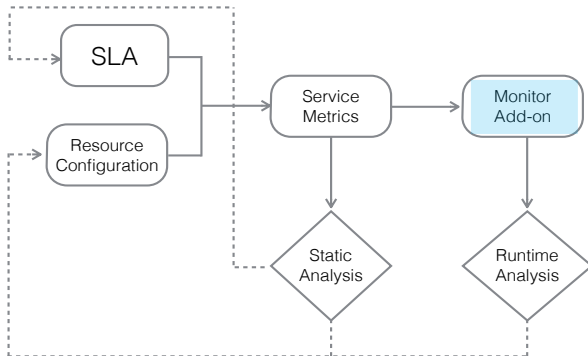
The Analysis Flow



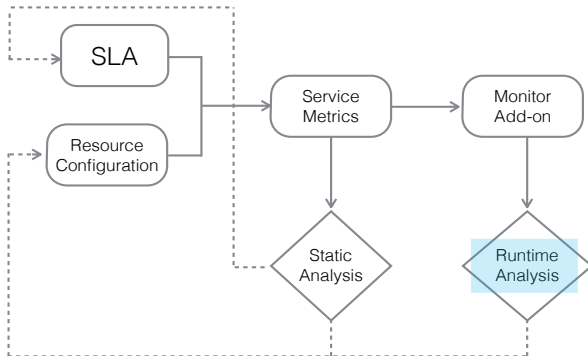
The Analysis Flow



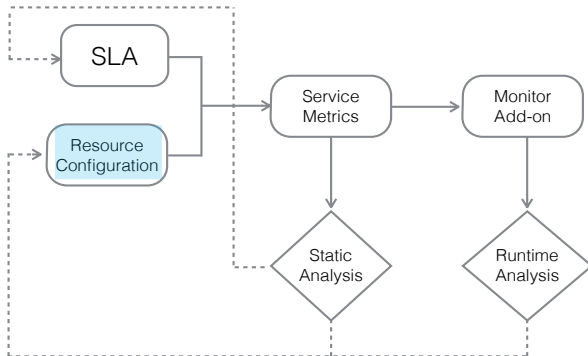
The Analysis Flow



The Analysis Flow



The Analysis Flow



SLAs and performance properties

Metrics for a Query API

"95% of requests is completed within 1 minute, 97% within 3 minutes and 98% within 5 minutes"

AVAILABILITY:

the **percentage of timely service provisioning requests**

"the service completes 6 queries per minute from 9:00 to 18:00 and 4 queries per minute otherwise"

CAPACITY:

the **service throughput**

"the service replies to a query request (with the result or with a failure) within 7 minutes"

RESPONSE TIME:

the **maximum response time**

Metrics for a Query API

"95% of requests is completed within 1 minute, 97% within 3 minutes and 98% within 5 minutes"

AVAILABILITY:

the **percentage of timely service provisioning requests**

"the service completes 6 queries per minute from 9:00 to 18:00 and 4 queries per minute otherwise"

CAPACITY:

the **service throughput**

"the service replies to a query request (with the result or with a failure) within 7 minutes"

RESPONSE TIME:

the **maximum response time**

Metrics for a Query API

"95% of requests is completed within 1 minute, 97% within 3 minutes and 98% within 5 minutes"

AVAILABILITY:

the **percentage of timely service provisioning requests**

"the service completes 6 queries per minute from 9:00 to 18:00 and 4 queries per minute otherwise"

CAPACITY:

the **service throughput**

"the service replies to a query request (with the result or with a failure) within 7 minutes"

RESPONSE TIME:

the **maximum response time**

Metrics for a Query API

"95% of requests is completed within 1 minute, 97% within 3 minutes and 98% within 5 minutes"

AVAILABILITY:

the **percentage of timely service provisioning requests**

"the service completes 6 queries per minute from 9:00 to 18:00 and 4 queries per minute otherwise"

CAPACITY:

the **service throughput**

"the service replies to a query request (with the result or with a failure) within 7 minutes"

RESPONSE TIME:

the **maximum response time**

¹ Cloud service level agreement standardisation guidelines (June 2014)

Other metrics ¹

security, data management, personal data protection

¹ Cloud service level agreement standardisation guidelines (June 2014)

SLA – Performance metrics

AVAILABILITY:

CAPACITY:

RESPONSE TIME:

Other metrics ¹

security, data management, personal data protection

¹ Cloud service level agreement standardisation guidelines (June 2014)

SLA – Performance metrics

AVAILABILITY:

the property of being accessible and usable upon demand

CAPACITY:

RESPONSE TIME:

Other metrics ¹

security, data management, personal data protection

¹ Cloud service level agreement standardisation guidelines (June 2014)

SLA – Performance metrics

AVAILABILITY:

the property of being **accessible and usable upon demand**

CAPACITY:

the **maximum amount of some property** of a cloud service.

RESPONSE TIME:

Other metrics ¹

security, data management, personal data protection

¹ Cloud service level agreement standardisation guidelines (June 2014)

SLA – Performance metrics

AVAILABILITY:

the property of being **accessible and usable upon demand**

CAPACITY:

the **maximum amount of some property** of a cloud service.

RESPONSE TIME:

the **time interval between** a cloud service **customer event** and a cloud service **provider response event**

Other metrics ¹

security, data management, personal data protection

¹ Cloud service level agreement standardisation guidelines (June 2014)

Metrics' formalisation: Service metric functions

- ▶ We formalise SLA metrics by means of **service metric functions**
- ▶ Assume $F(\tau)$ and $G(\tau, \delta)$ are two metric functions
- ▶ τ is a **time interval** of the form $[\mathsf{d}.t, \mathsf{d}'.t']$
 - d, d' are days ($\mathsf{d}, \mathsf{d}' \in \{1, \dots, 365\}$)
 - t, t' are seconds in the day ($t, t' \in \{0, \dots, 86399\}$)
- ▶ δ can be
 - an upper bound to the datasize (in bytes), or
 - a time bound for getting a reply, or
 - an upper bound to the number of resources used by the service.

Percentage of timely service provisioning requests (PTS)

FROM THE SLA

95% of requests is completed within 1 minute, 97% within 3 minutes and 98% within 5 minutes.

FORMALLY

$$PTS_s([1.0, 365.86399], x) = \begin{cases} 0,95 & \text{if } x = 60s \\ 0,97 & \text{if } x = 180s \\ 0,98 & \text{if } x = 300s \end{cases}$$

Service Throughput (ST)

FROM THE SLA

the service completes 6 queries per minute from 9:00 to 18:00 and 4 queries per minute otherwise.

FORMALLY

$$ST_s([1.t, 365.t'], 60) = \begin{cases} 4 & \text{if } t = 0 \text{ and } t' = 32399 \\ 6 & \text{if } t = 32400 \text{ and } t' = 64800 / 9am - 6pm \\ 4 & \text{if } t = 64801 \text{ and } t' = 86399 \end{cases}$$

Response Time (RT)

FROM THE SLA

the service replies to a query request (with the result or with a failure) within 7 minutes.

FORMALLY

$$RT_s([1.0, 365.86399]) = 420s$$

Static analysis

Behavioural types

- ▶ Abstract descriptions of programs
- ▶ Highlight the **information** for a particular property.
 - service time
 - number of virtual machines

APPROACH

- ① Extract behavioural types from a program
- ② Translate ① into abstract descriptions, e.g., equations
- ③ Solve ② to get, e.g., upper bound of service time

Behavioural types

- ▶ Abstract descriptions of programs
- ▶ Highlight the **information** for a particular property.
 - service time
 - number of virtual machines

APPROACH

- ① **Extract** behavioural types from a program
- ② **Translate** ① into abstract descriptions, e.g., equations
- ③ **Solve** ② to get, e.g., upper bound of service time

Behavioural types

- ▶ Abstract descriptions of programs
- ▶ Highlight the **information** for a particular property.
 - service time
 - number of virtual machines

APPROACH

- ① **Extract** behavioural types from a program
- ② **Translate** ① into abstract descriptions, e.g., equations
- ③ **Solve** ② to get, e.g., upper bound of service time

Behavioural types

- ▶ Abstract descriptions of programs
- ▶ Highlight the **information** for a particular property.
 - service time
 - number of virtual machines

APPROACH

- ① **Extract** behavioural types from a program
- ② **Translate** ① into abstract descriptions, e.g., equations
- ③ **Solve** ② to get, e.g., upper bound of service time

Example

```
String searchDB(Database DB,  
    String s)  
{  
    String u;  
    u = DB.query(s) ;  
    job(h) ;  
    .....  
    return u;  
}
```

- ▶ `job(h)` needs h CPU resources
- ▶ the time is $\frac{\text{needed CPU resources}}{\text{allocated CPU resources}}$

Example

Behavioural type:

```
String searchDB(Database DB,  
    String s)  
{  
    String u;  
    u = DB.query(s) ;  
    job(h) ;  
    .....  
    return u;  
}
```

```
Service.searchDB(a[x], b[y])  
{  
  
    DataBase.query(b[y]) ;  
    h/x ;  
} : _
```

Example

Behavioural type:

```
String searchDB(DataBase DB,
    String s)
{
    String u;
    u = DB.query(s) ;
    job(h) ;
    .....
    return u;
}

class DataBase {
    String query(String s) {
        .....
        job(k);
    }
}
```

```
Service.searchDB(a[x], b[y]) {

    DataBase.query(b[y]) ;
    h/x ;

} : _

DataBase.query(b[y]) {

    k/y

} : _
```

COST EQUATIONS

$$\text{searchDB}(x, y) = \text{query}(y) + h/x$$
$$\text{query}(y) = k/y$$

Example

Behavioural type:

```
String searchDB(DataBase DB,
    String s)
{
    String u;
    u = DB.query(s) ;
    job(h) ;
    .....
    return u;
}

class DataBase {
    String query(String s) {
        .....
        job(k);
    }
}
```

```
Service.searchDB(a[x], b[y]) {

    DataBase.query(b[y]) ;
    h/x ;

} : _

DataBase.query(b[y]) {

    k/y

} : _
```

COST EQUATIONS

$$\text{searchDB}(x, y) = \text{query}(y) + h/x$$
$$\text{query}(y) = k/y$$

Example

COST EQUATIONS

$\text{searchDB}(x, y) = \text{query}(y) + h/x$

$\text{query}(y) = k/y$

Example

COST EQUATIONS

$\text{searchDB}(x, y) = \text{query}(y) + h/x$

$\text{query}(y) = k/y$

Time estimation

$$\frac{h}{x} + \frac{k}{y}$$

Example

COST EQUATIONS

$$\text{searchDB}(x, y) = \text{query}(y) + h/x$$

$$\text{query}(y) = k/y$$

Time estimation

$$\frac{15}{x} + \frac{5}{y}$$

Example

COST EQUATIONS

$$\text{searchDB}(x, y) = \text{query}(y) + h/x$$

$$\text{query}(y) = k/y$$

Time estimation

$$\frac{15}{2} + \frac{5}{1}$$

Example

COST EQUATIONS

$\text{searchDB}(x, y) = \text{query}(y) + h/x$

$\text{query}(y) = k/y$

Time estimation

$$\frac{15}{2} + \frac{5}{1} = 12.5$$

Example

COST EQUATIONS

$$\text{searchDB}(x, y) = \text{query}(y) + h/x$$

$$\text{query}(y) = k/y$$

Time estimation

$$\frac{15}{2} + \frac{5}{1} = 12.5$$

$$ST_{\text{searchDB}}([1.t, 365.t'], 60) = \begin{cases} 4 & \text{if } t = 0 \text{ and } t' = 32399 \\ 6 & \text{if } t = 32400 \text{ and } t' = 64800 // 9am - 6pm \\ 4 & \text{if } t = 64801 \text{ and } t' = 86399 \end{cases}$$

$$RT_{\text{searchDB}}([1.0, 365.86399]) = 420s$$

Example

COST EQUATIONS

$$\text{searchDB}(x, y) = \text{query}(y) + h/x$$

$$\text{query}(y) = k/y$$

Time estimation

$$\frac{15}{2} + \frac{5}{1} = 12.5$$

$$ST_{\text{searchDB}}([1.t, 365.t'], 60) = \begin{cases} 4 & \text{if } t = 0 \text{ and } t' = 32399 \text{ 😊} \\ 6 & \text{if } t = 32400 \text{ and } t' = 64800 // 9am - 6pm \text{ 😞} \\ 4 & \text{if } t = 64801 \text{ and } t' = 86399 \text{ 😊} \end{cases}$$

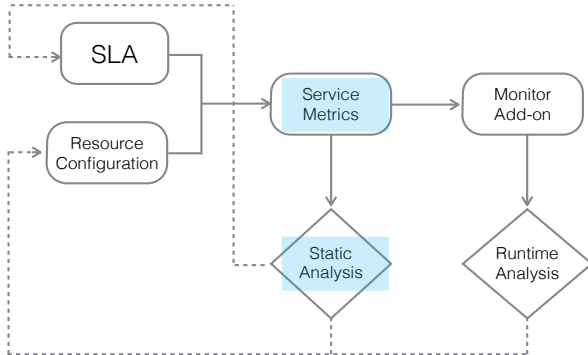
$$RT_{\text{searchDB}}([1.0, 365.86399]) = 420s \text{ 😊}$$

Break ST_{searchDB}

Adjustment

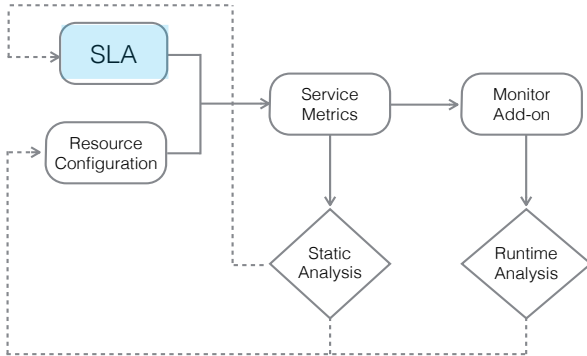
Time estimation

$$\frac{15}{2} + \frac{5}{1} = 12.5$$



Time estimation

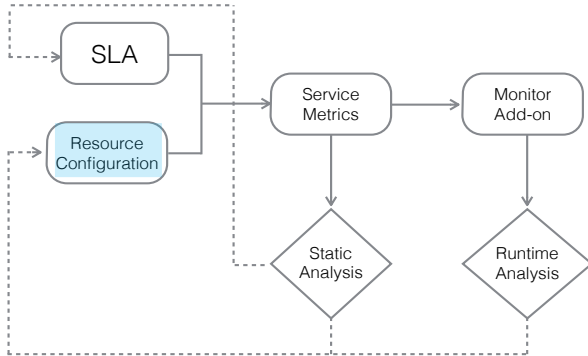
$$\frac{15}{2} + \frac{5}{1} = 12.5$$



Adjustment

Time estimation

$$\frac{15}{2} + \frac{5}{2} = 10$$



Time estimation

$$\frac{15}{2} + \frac{5}{2} = 10$$

$$ST_{\text{searchDB}}([1.t, 365.t'], 60) = \begin{cases} 4 & \text{if } t = 0 \text{ and } t' = 32399 \\ 6 & \text{if } t = 32400 \text{ and } t' = 64800 // 9am - 6pm \\ 4 & \text{if } t = 64801 \text{ and } t' = 86399 \end{cases}$$

$$RT_{\text{searchDB}}([1.0, 365.86399]) = 420s$$

Adjustment

Time estimation

$$\frac{15}{2} + \frac{5}{2} = 10$$

$$ST_{\text{searchDB}}([1.t, 365.t'], 60) = \begin{cases} 4 & \text{if } t = 0 \text{ and } t' = 32399 \text{ 😊} \\ 6 & \text{if } t = 32400 \text{ and } t' = 64800 // 9am - 6pm \text{ 😊} \\ 4 & \text{if } t = 64801 \text{ and } t' = 86399 \text{ 😊} \end{cases}$$

$$RT_{\text{searchDB}}([1.0, 365.86399]) = 420s \text{ 😊}$$

Satisfy both

Run-time analysis

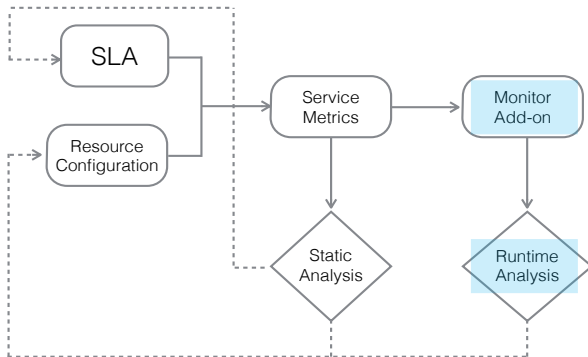
Run-time monitoring

- ▶ Observation:
- ▶ Reaction:

Monitoring metrics

Run-time monitoring

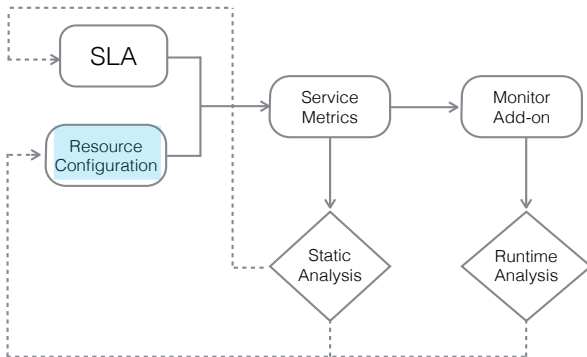
- **Observation**: measurements on services are taken
- **Reaction**:



Monitoring metrics

Run-time monitoring

- ▶ **Observation**: measurements on services are taken
- ▶ **Reaction**: adjustments for resource allocations



Service time monitor

- ▶ Intercepts all the HTTP invocations/replies to/from a service
- ▶ Records the time taken by every request to be completed

```
void monitor_service_time() {  
    (service,method,msg,client,m_id) = HttpRequest.intercept();  
    time_start = time();  
    reply = service.method(msg);  
    time_end = time();  
    HttpResponse.send(client,reply,m_id);  
    log(m_id,time_start,time_end);  
}
```

Percentage of timely service provisioning requests

percentage method

takes as input a time window and returns `true` if the percentage of requests complies with the definition of PTS_s

```
boolean percentage(Log_file d, Time t_begin, Time t_end){
    boolean v = true ;

    /* retrieve from the log the total number of messages served in the time window */
    nmb_msg = get_total_messages(d, t_begin, t_end) ;

    /* check whether the SLA percentages of served requests
       correspond to the observed ones */
    nmb_msg_completed = find(t_begin, t_end, 60) ;
    v = v && (nmb_msg_completed/nmb_msg <= 0.95) ; //95% in 1 min
    nmb_msg_completed = find(t_begin, t_end, 180) ;
    v = v && (nmb_msg_completed/nmb_msg <= 0.97) ; //97% in 3 mins
    nmb_msg_completed = find(t_begin, t_end, 300) ;
    v = v && (nmb_msg_completed/nmb_msg <= 0.98) ; //98% in 5 mins

    return v;
}
```

Service Throughput

throughput method

takes as input a logfile and a time window and returns `true` if the throughput complies with the definition of ST_s

```
boolean throughput(Log_file d, Time t_begin, Time t_end){
    int daily = 0;
    int nightly = 0;

    /* collects the number of the served requests during the two
       specified time-frames */
    for each (m_id, time_init, time_end) in d {
        if ((time_init >= 32400) && (time_init <= 64800)) // 9:00–18:00
            daily = daily + 1 ;
        else nightly = nightly + 1 ;
    }

    /* return true if 6 queries per minute are completed in 9:00-18:00
       and 4 queries per minute in the remaining time */
    return ( ((daily/60*9)>6) && ((nightly/60*15)>4) );
}
```

Example

$$ST_{\text{searchDB}}([1.t, 365.t'], 60) = \begin{cases} 4 & \text{if } t = 0 \text{ and } t' = 32399 \\ 6 & \text{if } t = 32400 \text{ and } t' = 64800 // 9am - 6pm \\ 4 & \text{if } t = 64801 \text{ and } t' = 86399 \end{cases}$$

$$RT_{\text{searchDB}}([1.0, 365.86399]) = 420s$$

- ▶ [Static Analysis] The initial 2 CPU resources respect the SLA. 😊
- ▶ The `throughput` method records only 4 requests per minute. 😞

Example

$$ST_{\text{searchDB}}([1.t, 365.t'], 60) = \begin{cases} 4 & \text{if } t = 0 \text{ and } t' = 32399 \\ 6 & \text{if } t = 32400 \text{ and } t' = 64800 // 9am - 6pm \\ 4 & \text{if } t = 64801 \text{ and } t' = 86399 \end{cases}$$

$$RT_{\text{searchDB}}([1.0, 365.86399]) = 420s$$

- ▶ [Static Analysis] The initial 2 CPU resources respect the SLA. 😊
- ▶ The `throughput` method records only 4 requests per minute. 😞

Example

$$ST_{\text{searchDB}}([1.t, 365.t'], 60) = \begin{cases} 4 & \text{if } t = 0 \text{ and } t' = 32399 \\ 6 & \text{if } t = 32400 \text{ and } t' = 64800 // 9am - 6pm \\ 4 & \text{if } t = 64801 \text{ and } t' = 86399 \end{cases}$$

$$RT_{\text{searchDB}}([1.0, 365.86399]) = 420s$$

- ▶ [Static Analysis] The initial 2 CPU resources respect the SLA. 😊
- ▶ The `throughput` method records only 4 requests per minute. 😞

Example

$$ST_{\text{searchDB}}([1.t, 365.t'], 60) = \begin{cases} 4 & \text{if } t = 0 \text{ and } t' = 32399 \\ 6 & \text{if } t = 32400 \text{ and } t' = 64800 // 9am - 6pm \\ 4 & \text{if } t = 64801 \text{ and } t' = 86399 \end{cases}$$

$$RT_{\text{searchDB}}([1.0, 365.86399]) = 420s$$

- ▶ [Static Analysis] The initial 2 CPU resources respect the SLA. 😊
- ▶ The `throughput` method records only 4 requests per minute. 😞

Reaction

Configure resource allocation

reaction method

verifies every 300s whether the percentage of timely service provisioning requests is reached and, in case of failures, adds one more CPU

```
void reaction(Service s) {  
    Time t ; Bool v ;  
    t = time() ;  
    idle(300) ;  
    v = percentage(d,t, t+300) ;  
    if (!v) MonitoringPlatform ! allocate(s) ;  
}
```

Example

$$ST_{\text{searchDB}}([1.t, 365.t'], 60) = \begin{cases} 4 & \text{if } t = 0 \text{ and } t' = 32399 \\ 6 & \text{if } t = 32400 \text{ and } t' = 64800 // 9am - 6pm \\ 4 & \text{if } t = 64801 \text{ and } t' = 86399 \end{cases}$$

$$RT_{\text{searchDB}}([1.0, 365.86399]) = 420s$$

- ▶ [Static Analysis] The initial 2 CPU resources respect the SLA. 😊
- ▶ The `throughput` method records only 4 requests per minute. 😞
- ▶ The `reaction` method requests 1 additional CPU resource. 😊
- ▶ Waste of money and resources during the night. 😞

Example

$$ST_{\text{searchDB}}([1.t, 365.t'], 60) = \begin{cases} 4 & \text{if } t = 0 \text{ and } t' = 32399 \\ 6 & \text{if } t = 32400 \text{ and } t' = 64800 // 9am - 6pm \\ 4 & \text{if } t = 64801 \text{ and } t' = 86399 \end{cases}$$

$$RT_{\text{searchDB}}([1.0, 365.86399]) = 420s$$

- ▶ [Static Analysis] The initial 2 CPU resources respect the SLA. 😊
- ▶ The `throughput` method records only 4 requests per minute. 😞
- ▶ The `reaction` method requests 1 additional CPU resource. 😊
- ▶ Waste of money and resources during the night. 😞

Example

$$ST_{\text{searchDB}}([1.t, 365.t'], 60) = \begin{cases} 4 & \text{if } t = 0 \text{ and } t' = 32399 \\ 6 & \text{if } t = 32400 \text{ and } t' = 64800 // 9am - 6pm \\ 4 & \text{if } t = 64801 \text{ and } t' = 86399 \end{cases}$$

$$RT_{\text{searchDB}}([1.0, 365.86399]) = 420s$$

- ▶ [Static Analysis] The initial 2 CPU resources respect the SLA. 😊
- ▶ The `throughput` method records only 4 requests per minute. 😞
- ▶ The `reaction` method requests 1 additional CPU resource. 😊
- ▶ Waste of money and resources during the night. 😞

Example

Budget metric function

$$\text{Budget}_{\text{searchDB}}([1.t, 365.t']) = \begin{cases} 40 & \text{if } t = 0 \text{ and } t' = 32399 \\ 60 & \text{if } t = 32400 \text{ and } t' = 64800 // 9am - 6pm \\ 40 & \text{if } t = 64801 \text{ and } t' = 86399 \end{cases}$$

- ▶ Each CPU resource costs 10
- ▶ The static analysis will approve $\text{Budget}_{\text{searchDB}}$. 😊
- ▶ When the runtime CPU reallocation is triggered by the throughput monitor, the nightly budget is not met anymore. ☹️
- ▶ The `budget_monitor` reacts by requiring a deallocation of half of the CPU units during the night. 😊

Example

Budget metric function

$$\text{Budget}_{\text{searchDB}}([1.t, 365.t']) = \begin{cases} 40 & \text{if } t = 0 \text{ and } t' = 32399 \\ 60 & \text{if } t = 32400 \text{ and } t' = 64800 // 9am - 6pm \\ 40 & \text{if } t = 64801 \text{ and } t' = 86399 \end{cases}$$

- ▶ Each CPU resource costs 10
- ▶ The static analysis will approve $\text{Budget}_{\text{searchDB}}$. 😊
- ▶ When the runtime CPU reallocation is triggered by the throughput monitor, the nightly budget is not met anymore. 😞
- ▶ The `budget_monitor` reacts by requiring a deallocation of half of the CPU units during the night. 😊

Example

Budget metric function

$$\text{Budget}_{\text{searchDB}}([1.t, 365.t']) = \begin{cases} 40 & \text{if } t = 0 \text{ and } t' = 32399 \\ 60 & \text{if } t = 32400 \text{ and } t' = 64800 // 9am - 6pm \\ 40 & \text{if } t = 64801 \text{ and } t' = 86399 \end{cases}$$

- ▶ Each CPU resource costs 10
- ▶ The static analysis will approve $\text{Budget}_{\text{searchDB}}$. 😊
- ▶ When the runtime CPU reallocation is triggered by the throughput monitor, the nightly budget is not met anymore. 😞
- ▶ The `budget_monitor` reacts by requiring a deallocation of half of the CPU units during the night. 😊

Example

Budget metric function

$$\text{Budget}_{\text{searchDB}}([1.t, 365.t']) = \begin{cases} 40 & \text{if } t = 0 \text{ and } t' = 32399 \\ 60 & \text{if } t = 32400 \text{ and } t' = 64800 // 9am - 6pm \\ 40 & \text{if } t = 64801 \text{ and } t' = 86399 \end{cases}$$

- ▶ Each CPU resource costs 10
- ▶ The static analysis will approve $\text{Budget}_{\text{searchDB}}$. 😊
- ▶ When the runtime CPU reallocation is triggered by the throughput monitor, the nightly budget is not met anymore. 😞
- ▶ The `budget_monitor` reacts by requiring a deallocation of half of the CPU units during the night. 😊

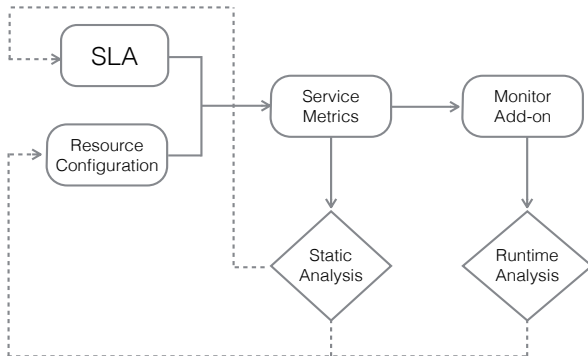
Example

Budget metric function

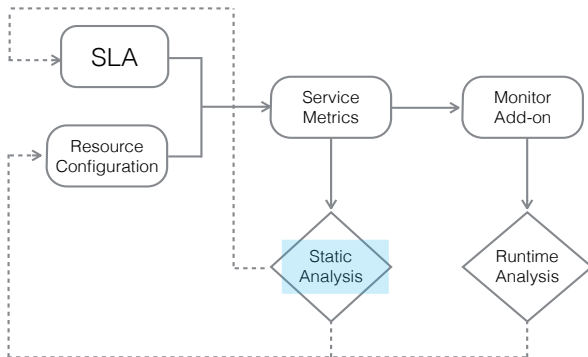
$$\text{Budget}_{\text{searchDB}}([1.t, 365.t']) = \begin{cases} 40 & \text{if } t = 0 \text{ and } t' = 32399 \\ 60 & \text{if } t = 32400 \text{ and } t' = 64800 // 9am - 6pm \\ 40 & \text{if } t = 64801 \text{ and } t' = 86399 \end{cases}$$

- ▶ Each CPU resource costs 10
- ▶ The static analysis will approve $\text{Budget}_{\text{searchDB}}$. 😊
- ▶ When the runtime CPU reallocation is triggered by the throughput monitor, the nightly budget is not met anymore. 😞
- ▶ The `budget_monitor` reacts by requiring a deallocation of half of the CPU units during the night. 😊

Summary

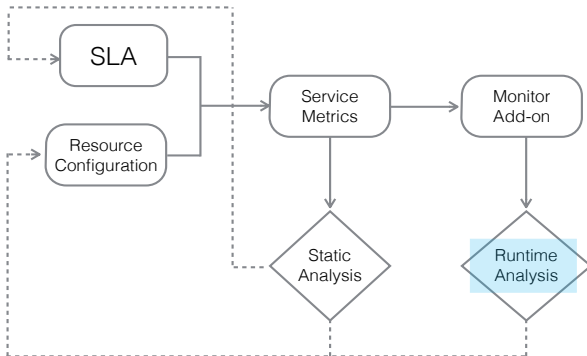


Summary



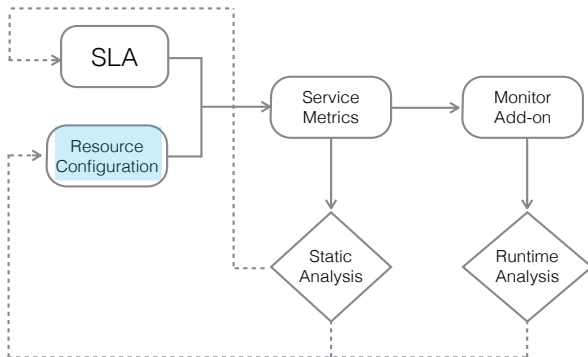
- ▶ Maximum number of virtual machines
- ▶ Upper bound of response time

Summary



- Availability
- Budget compliances

Summary



- A monitoring platform for resource (de-)allocation.

Summary

Static Analysis

- [Static analysis of cloud elasticity](#). (PPDP'15).
A. Garcia, C. Laneve, M. Lienhardt
- [Time Complexity of Concurrent Programs](#). (FACS'15).
E. Giachino, E. B. Johnsen, C. Laneve, K. I Pun

Runtime monitoring

- [Formal verification of service level agreements through distributed monitoring](#). (ESOCC'15).
B. Nobakht, S. de Gouw, and F. S. de Boer

SLA metric approach

- [Statically and Dynamically verifiable SLA metrics](#).
(Theory and Practice of Formal Methods).
E. Giachino, S. De Gouw, C. Laneve, and B. Nobakht.