



NodeCompass: A new Paradigm for Performance Evaluation in Event-based Node.js

Progress Report Meeting

Hervé KABAMBA

PhD Candidate

Supervisor: Michel Dagenais

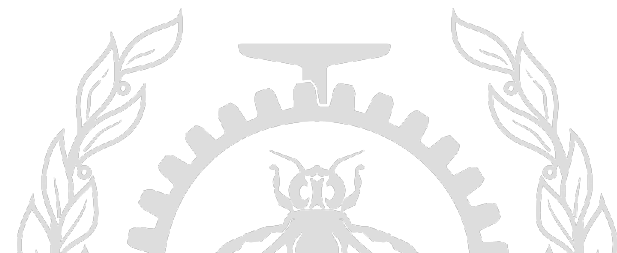
May 16, 2022

Polytechnique Montréal

Département de Génie Informatique et Génie Logiciel

Agenda

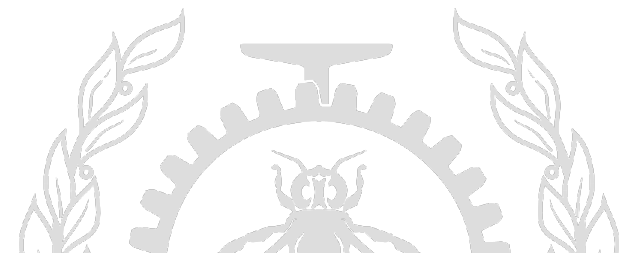
1. Introduction



Agenda

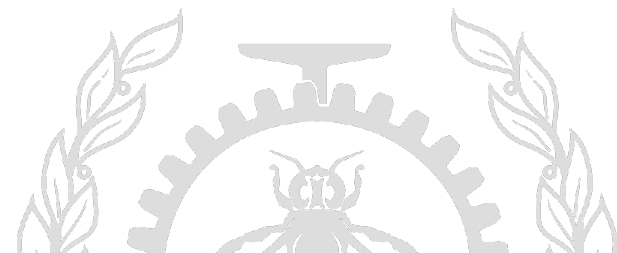
1. Introduction

2. Architecture



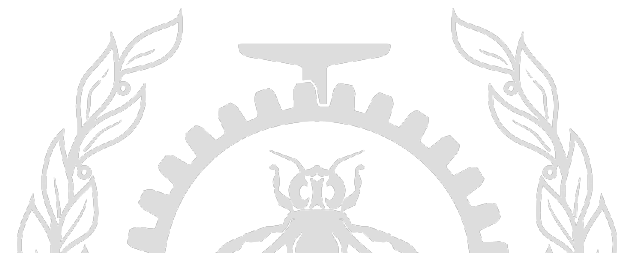
Agenda

1. Introduction
2. Architecture
3. Use Cases



Agenda

1. Introduction
2. Architecture
3. Use Cases
4. Bibliography



Introduction

NodeCompass: A new Paradigm?

Introduction

NodeCompass: A new Paradigm?

Yes. Performance Analysis is done differently

Introduction

NodeCompass: A new Paradigm?

Yes. Performance Analysis is done differently

We consider all layers interacting in the system as part of the Perf. Analysis

Introduction

NodeCompass: A new Paradigm?

Yes. Performance Analysis is done differently

We consider all layers interacting in the system as part of the Perf. Analysis

Reduced Instrumentation overhead by using LTTng in the JavaScript land

Introduction

NodeCompass: A new Paradigm?

Yes. Performance Analysis is done differently

We consider all layers interacting in the system as part of the Perf. Analysis

Reduced Instrumentation overhead by using LTTng in the JavaScript land

Focus on High Level of Granularity in the Perf. Analysis

Introduction

NodeCompass: A new Paradigm?

Yes. Performance Analysis is done differently

We consider all layers interacting in the system as part of the Perf. Analysis

Reduced Instrumentation overhead by using LTTng in the JavaScript land

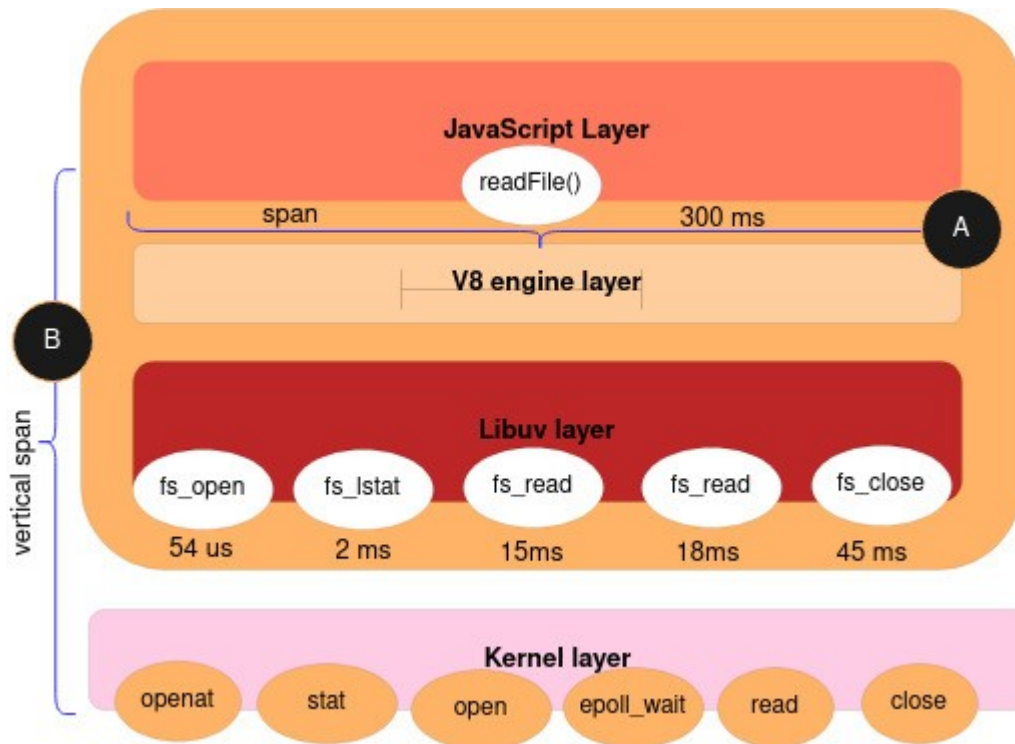
Focus on High Level of Granularity in the Perf. Analysis

Uncover Bugs, Bottlenecks, Root causes, Race conditions either in Node.js internals or in JavaScript land

Introduction

Why is NodeCompass different?

Current approaches :

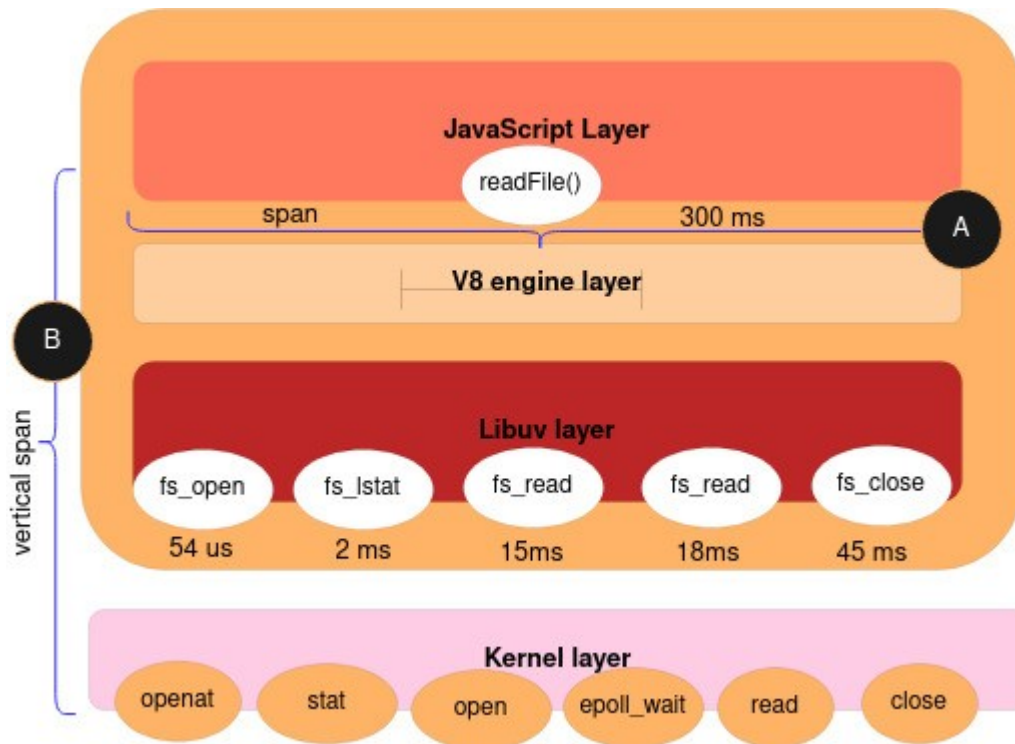


Introduction

Why is NodeCompass different?

Current approaches :

- Focus on higher layer

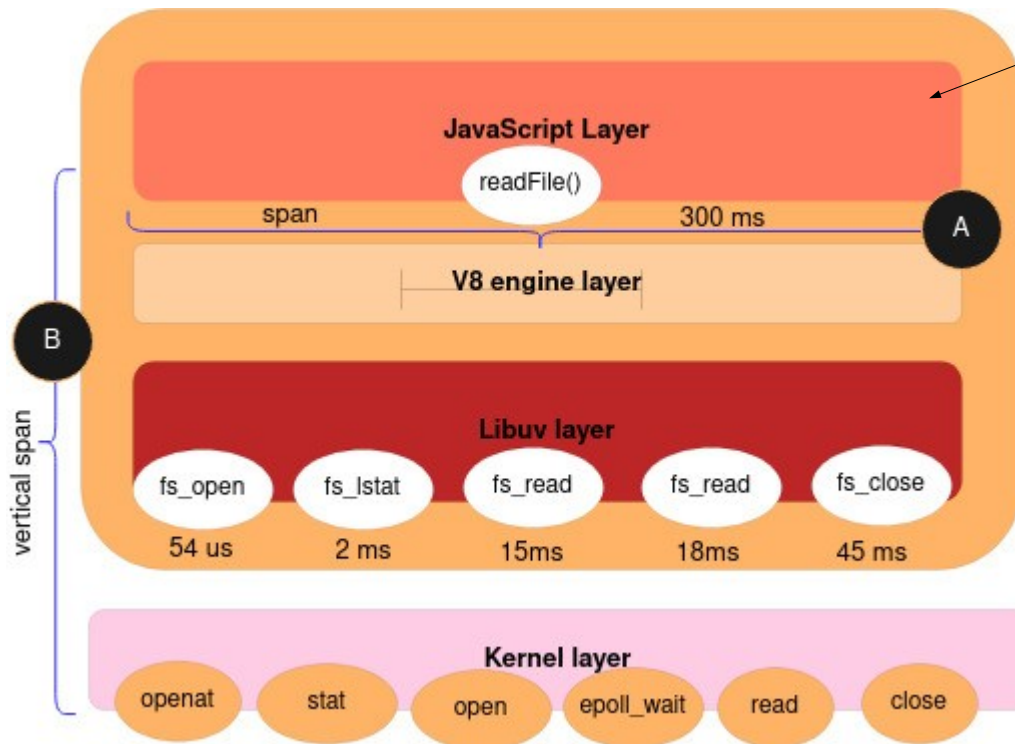


Introduction

Why is NodeCompass different?

Current approaches :

- Focus on higher layer

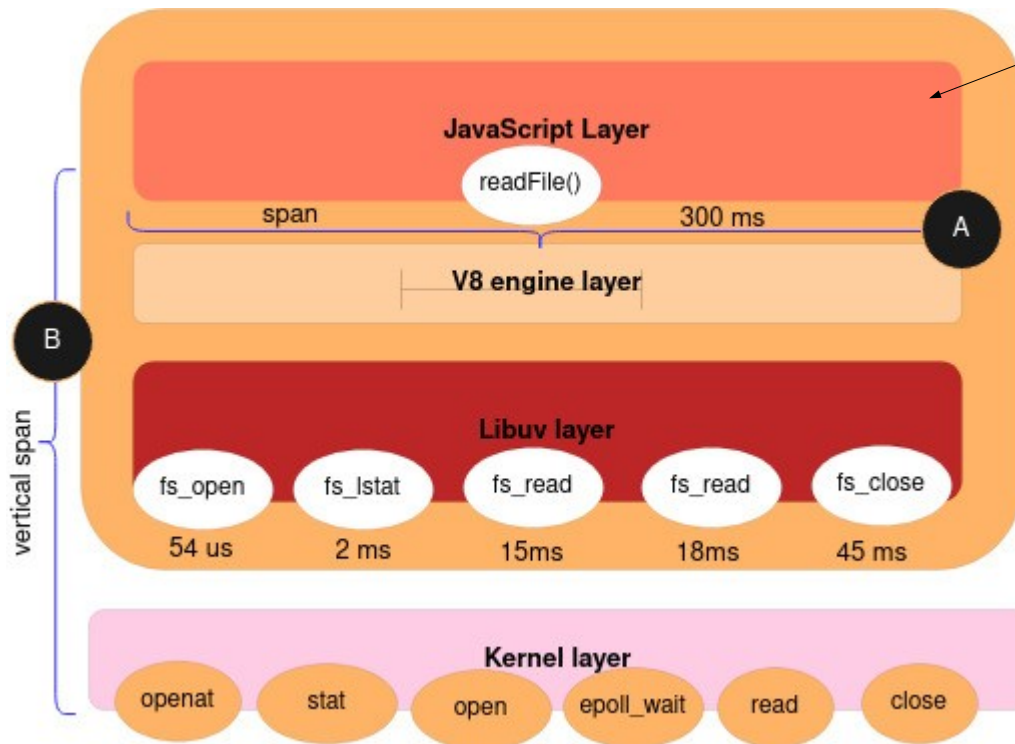


Introduction

Why is NodeCompass different?

Current approaches :

- Focus on higher layer
- Distributed tracers

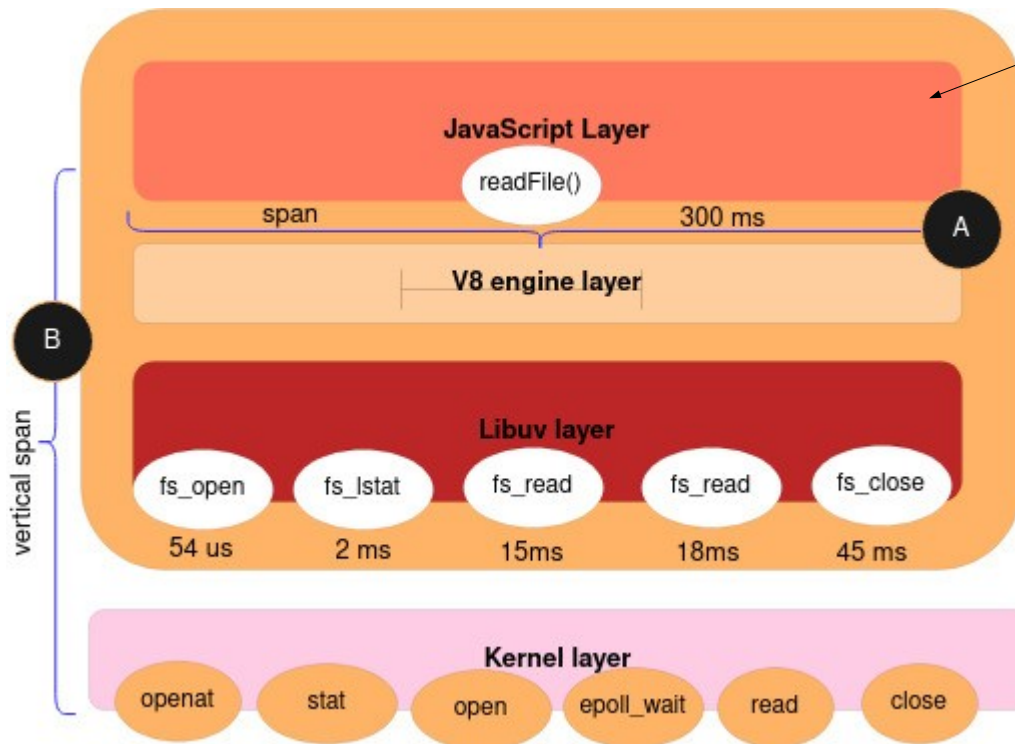


Introduction

Why is NodeCompass different?

Current approaches :

- Focus on higher layer
- Distributed tracers
- Spans are exposed

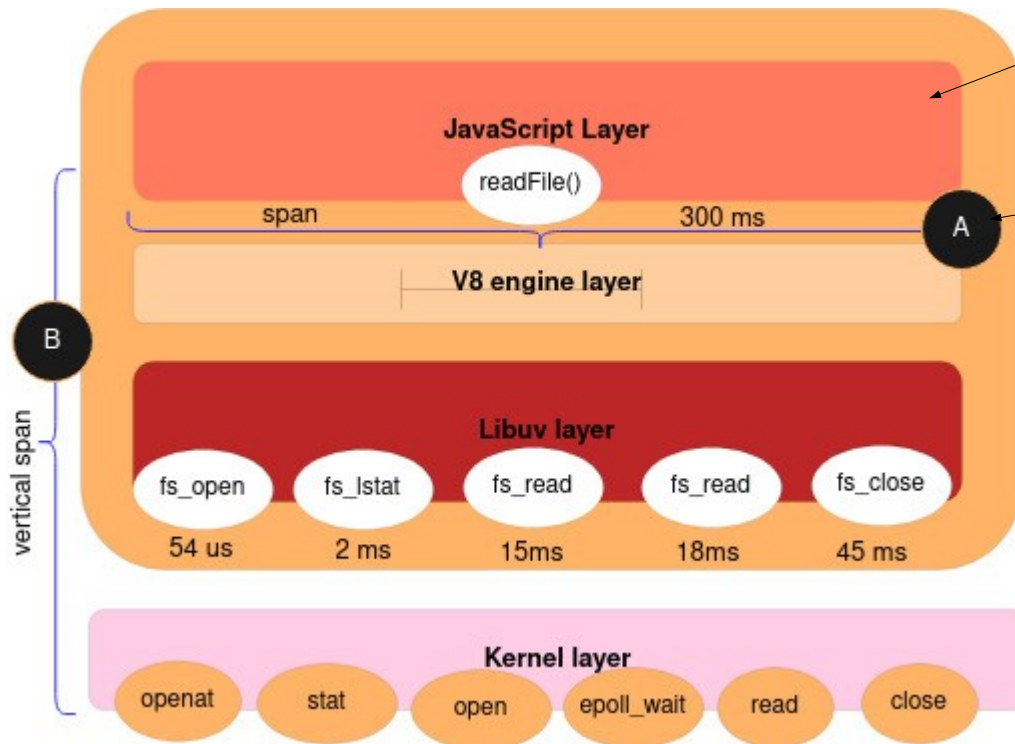


Introduction

Why is NodeCompass different?

Current approaches :

- Focus on higher layer
- Distributed tracers
- Spans are exposed

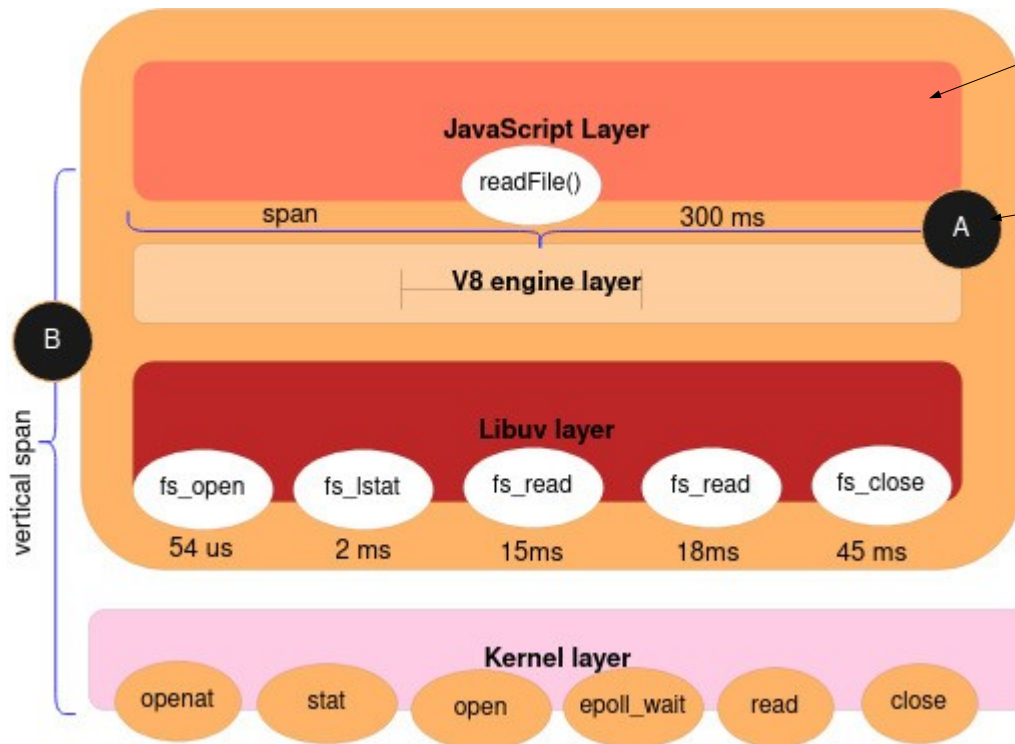


Introduction

Why is NodeCompass different?

Current approaches :

- Focus on higher layer
- Distributed tracers
- Spans are exposed



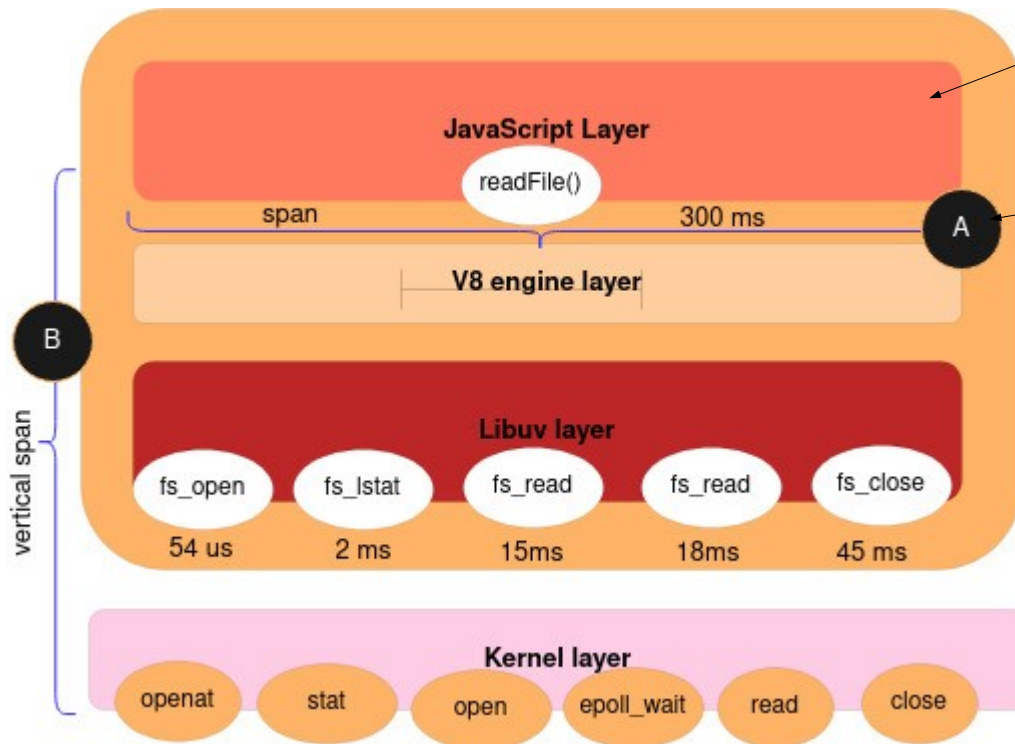
Cons

Introduction

Why is NodeCompass different?

Current approaches :

- Focus on higher layer
- Distributed tracers
- Spans are exposed



Cons

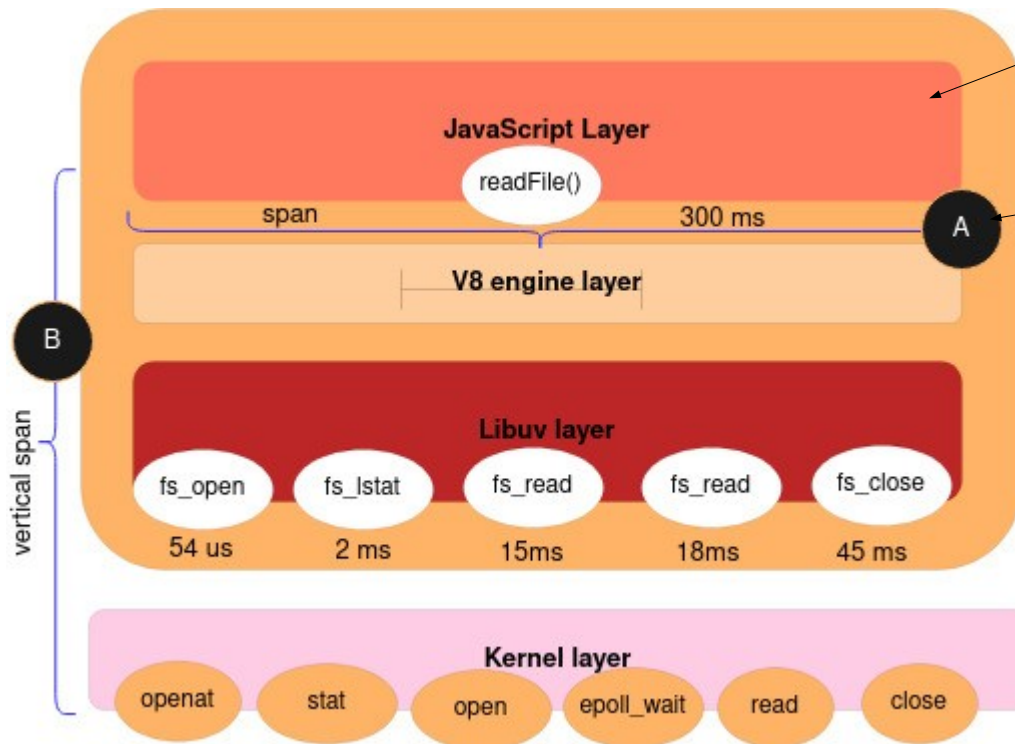
- Give only indications of the function latency

Introduction

Why is NodeCompass different?

Current approaches :

- Focus on higher layer
- Distributed tracers
- Spans are exposed



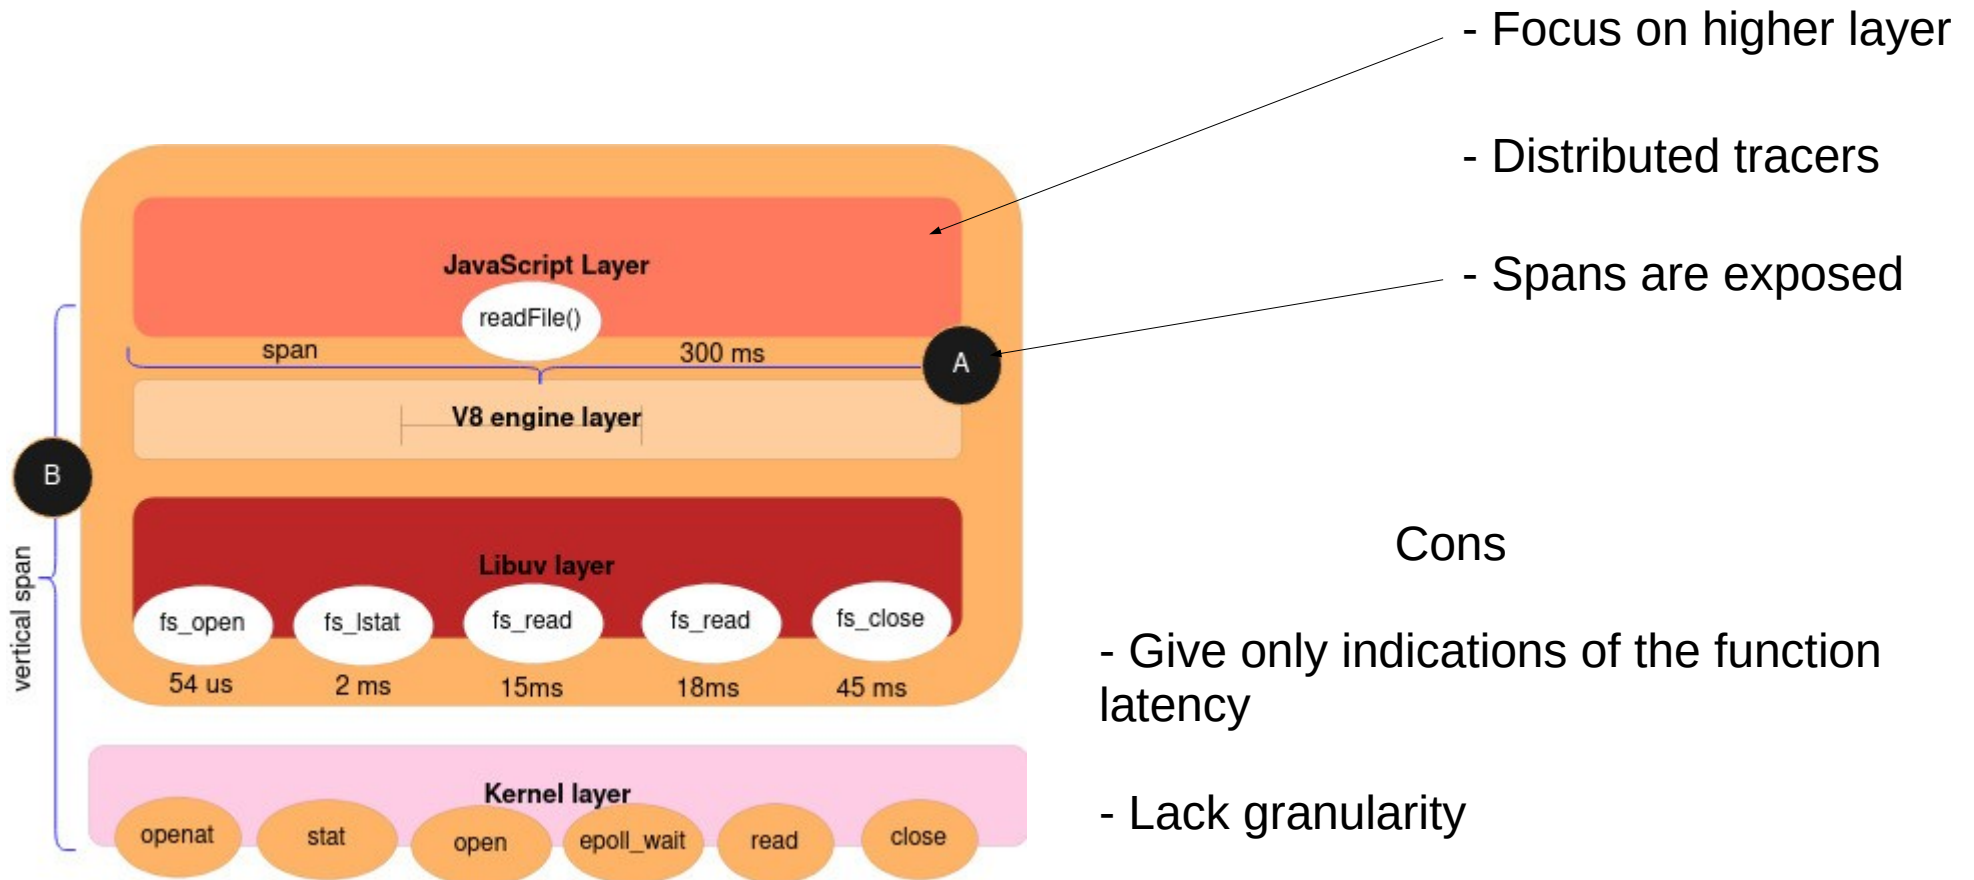
Cons

- Give only indications of the function latency
- Lack granularity

Introduction

Why is NodeCompass different?

Current approaches :

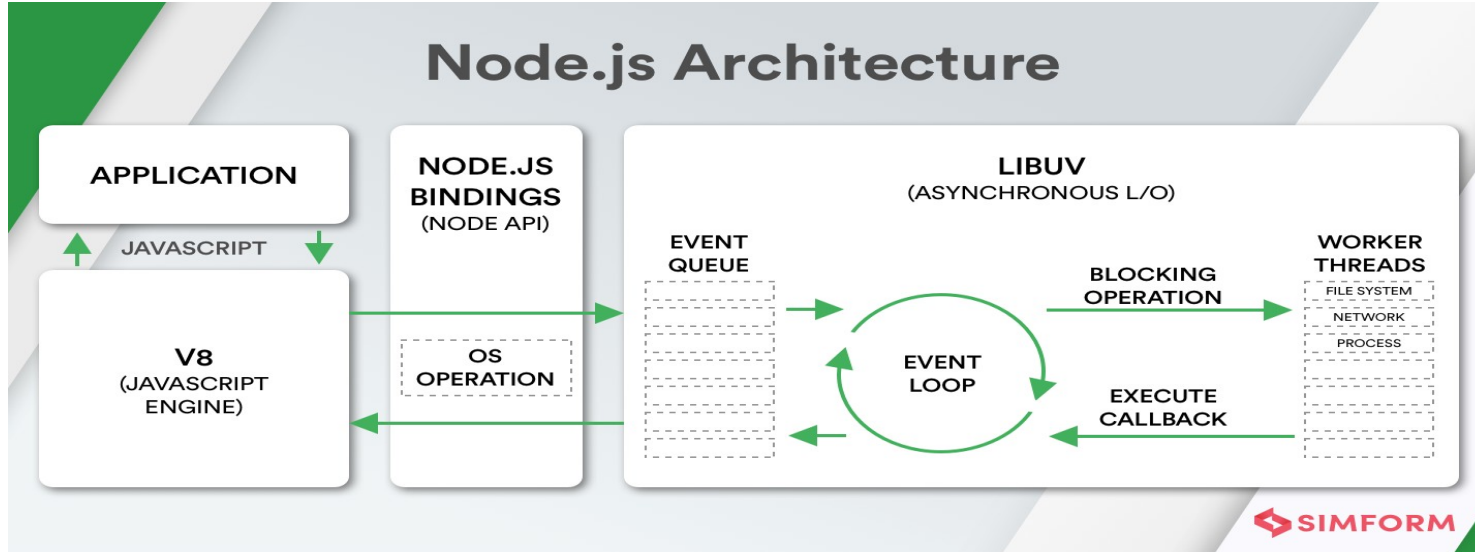


Cons

- Give only indications of the function latency
- Lack granularity
- Too much effort to understand the root cause based on the given information

Introduction

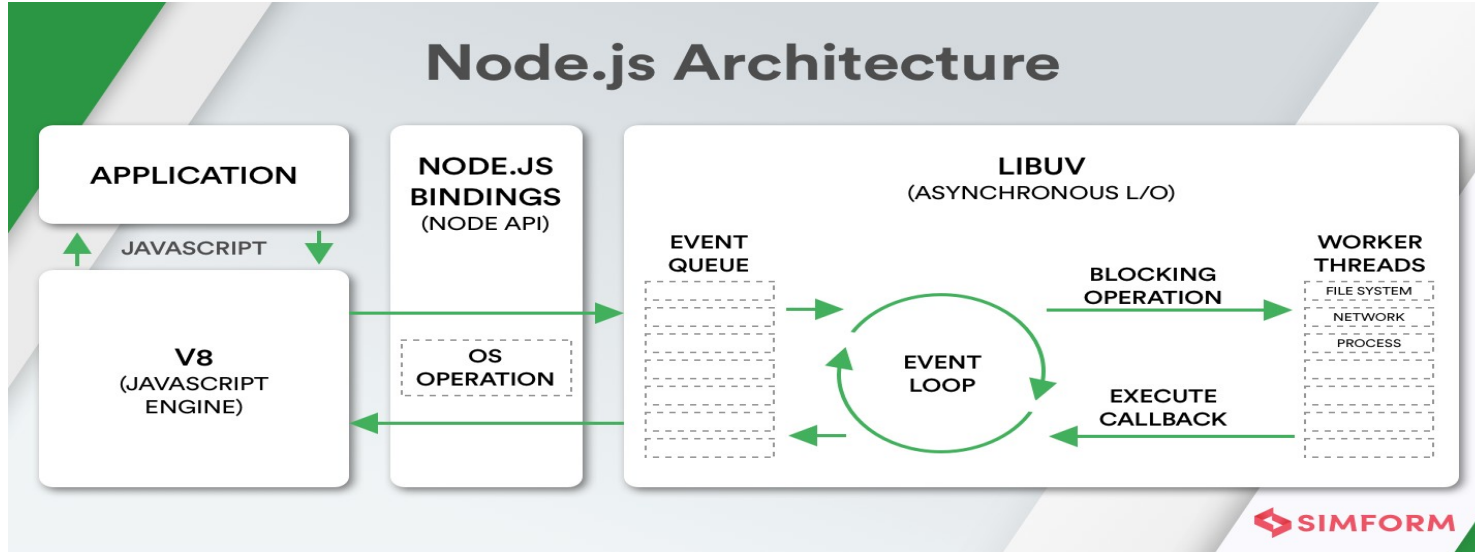
NodeCompass approach for Perf. Analysis



Introduction

NodeCompass approach for Perf. Analysis

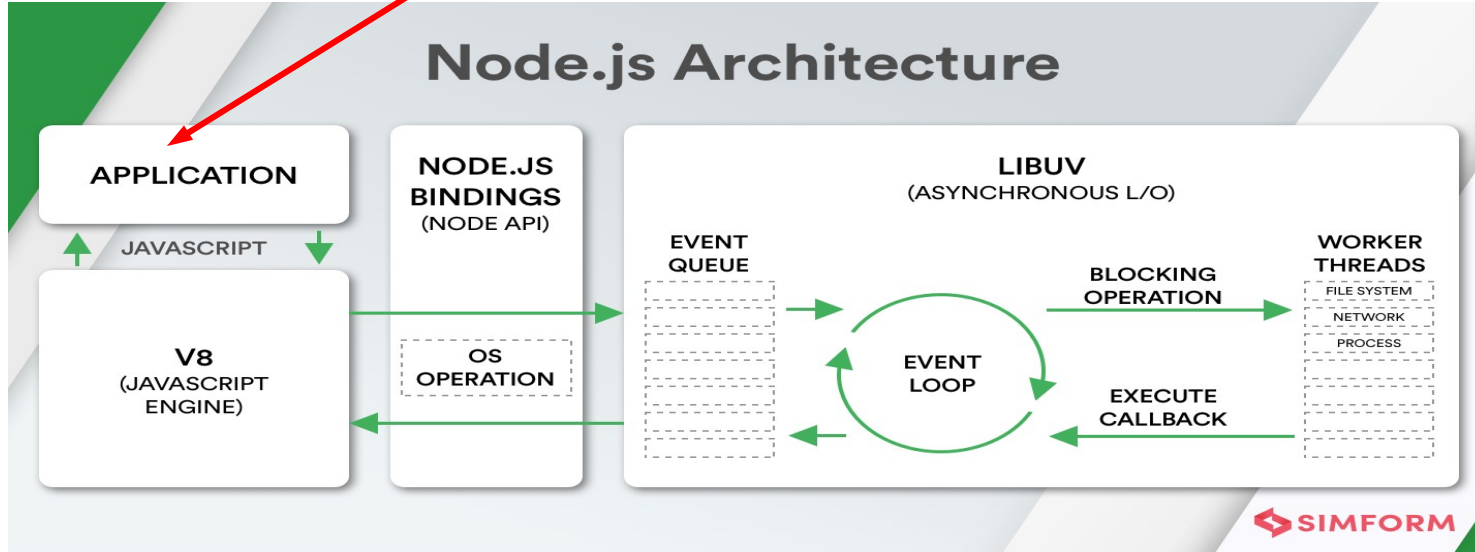
Tracepoints are inserted in the JavaScript land



Introduction

NodeCompass approach for Perf. Analysis

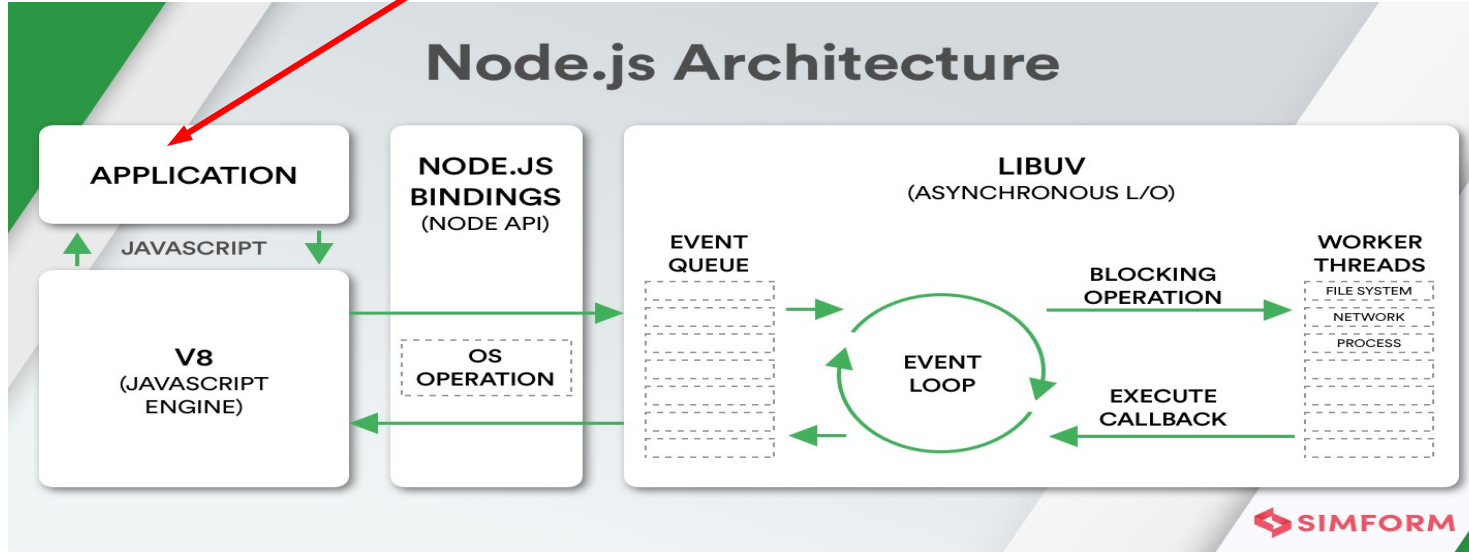
Tracepoints are inserted in the JavaScript land



Introduction

NodeCompass approach for Perf. Analysis

Tracepoints are inserted in the JavaScript land

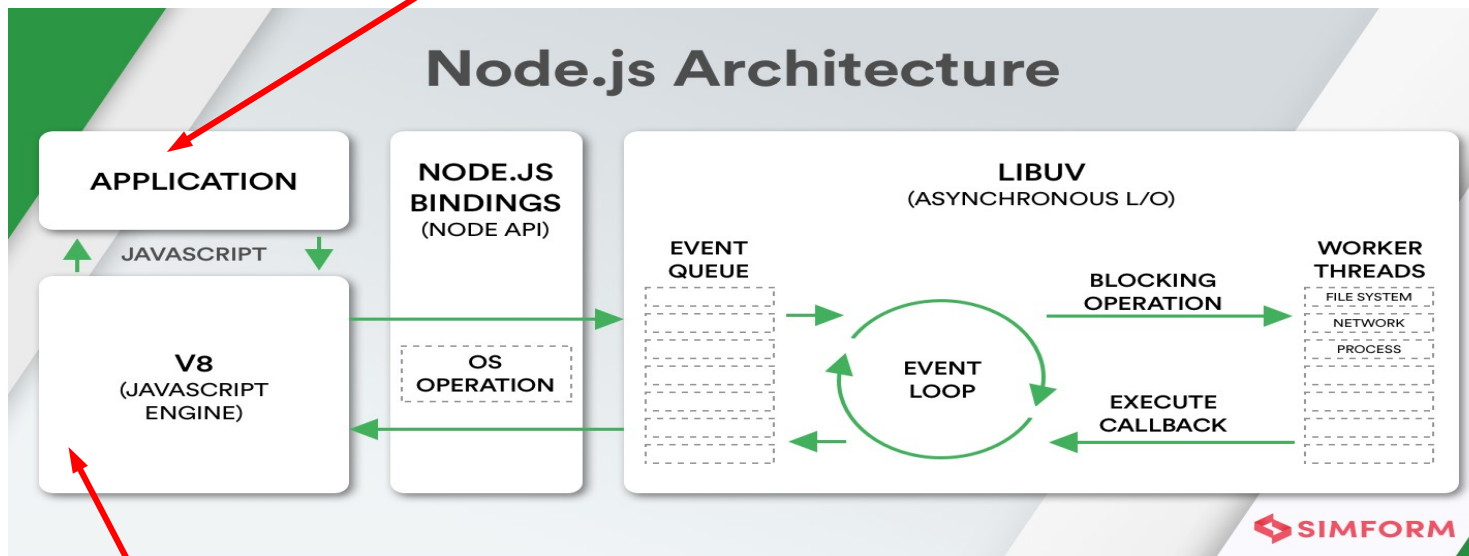


Tracepoints are inserted in the V8 OS

Introduction

NodeCompass approach for Perf. Analysis

Tracepoints are inserted in the JavaScript land

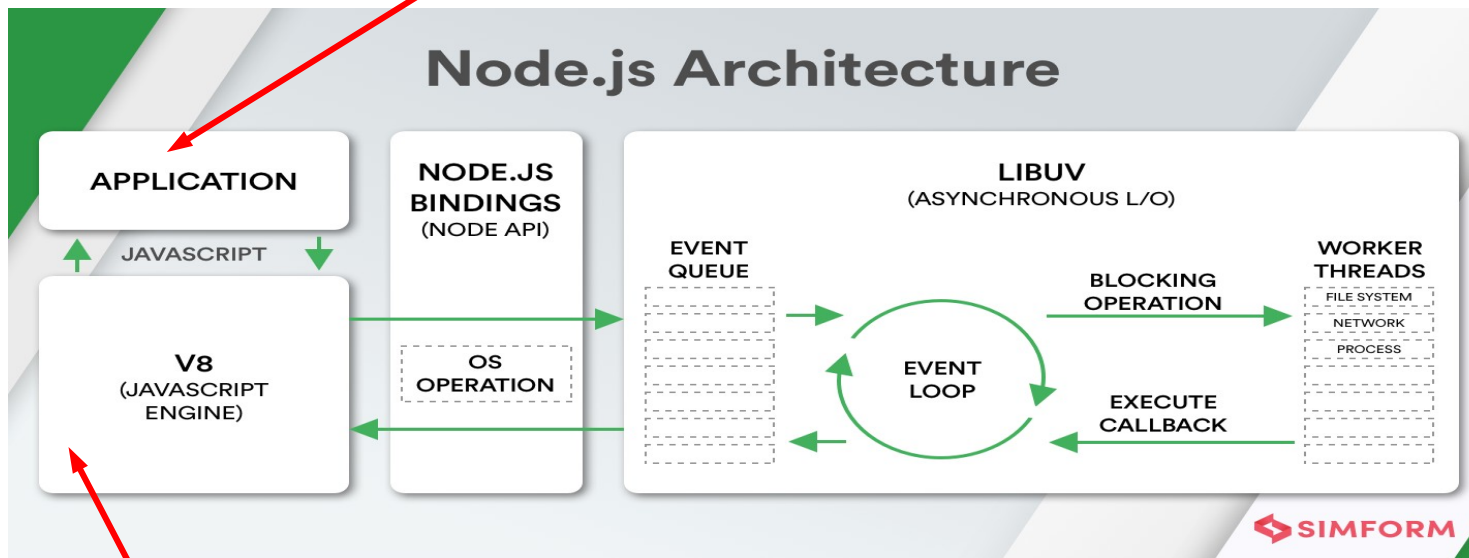


Tracepoints are inserted in the V8 OS

Introduction

NodeCompass approach for Perf. Analysis

Tracepoints are inserted in the JavaScript land



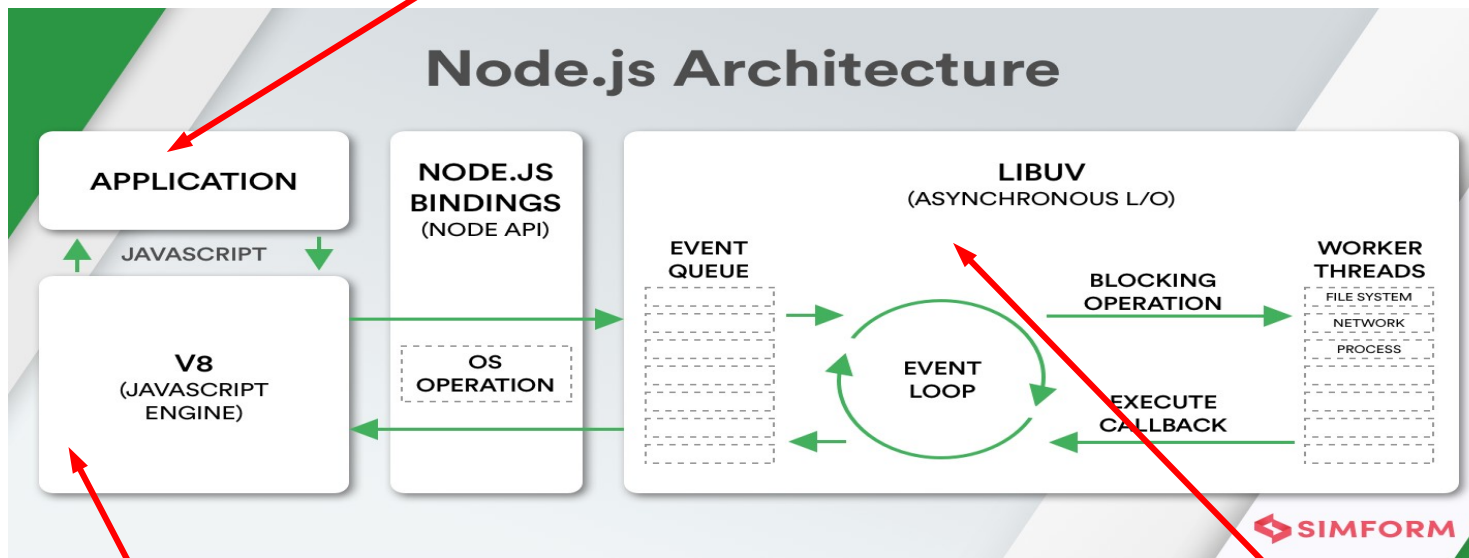
Tracepoints are inserted in the orchestration engine

Tracepoints are inserted in the V8 OS

Introduction

NodeCompass approach for Perf. Analysis

Tracepoints are inserted in the JavaScript land



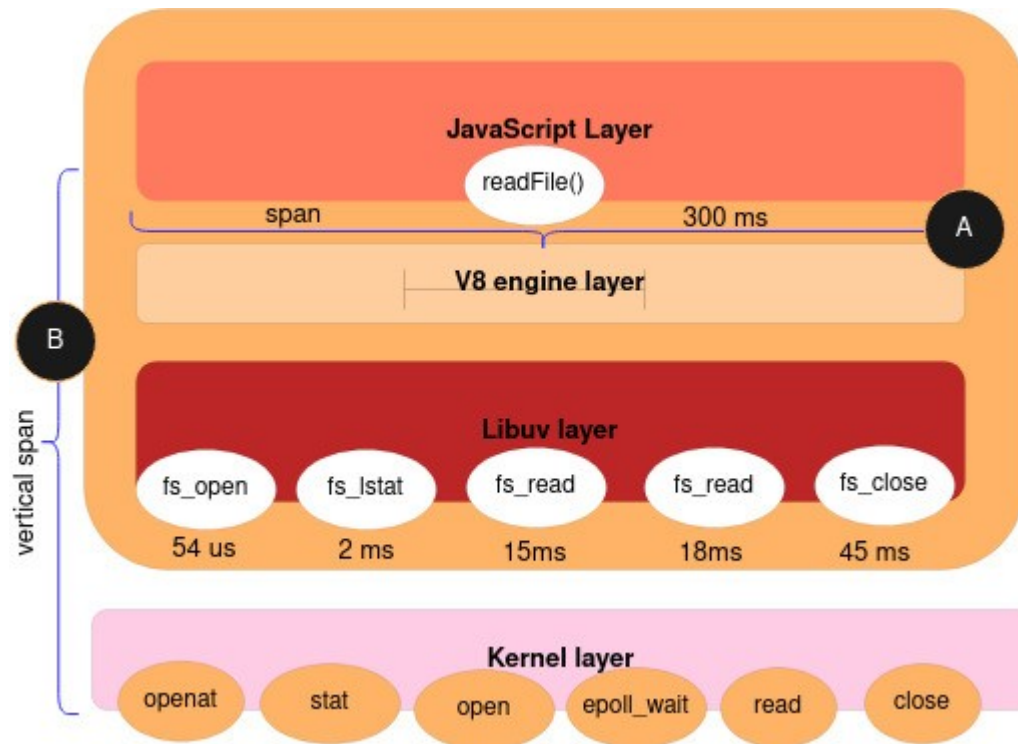
Tracepoints are inserted in the orchestration engine

Tracepoints are inserted in the V8 OS

Introduction

NodeCompass approach

New paradigm:

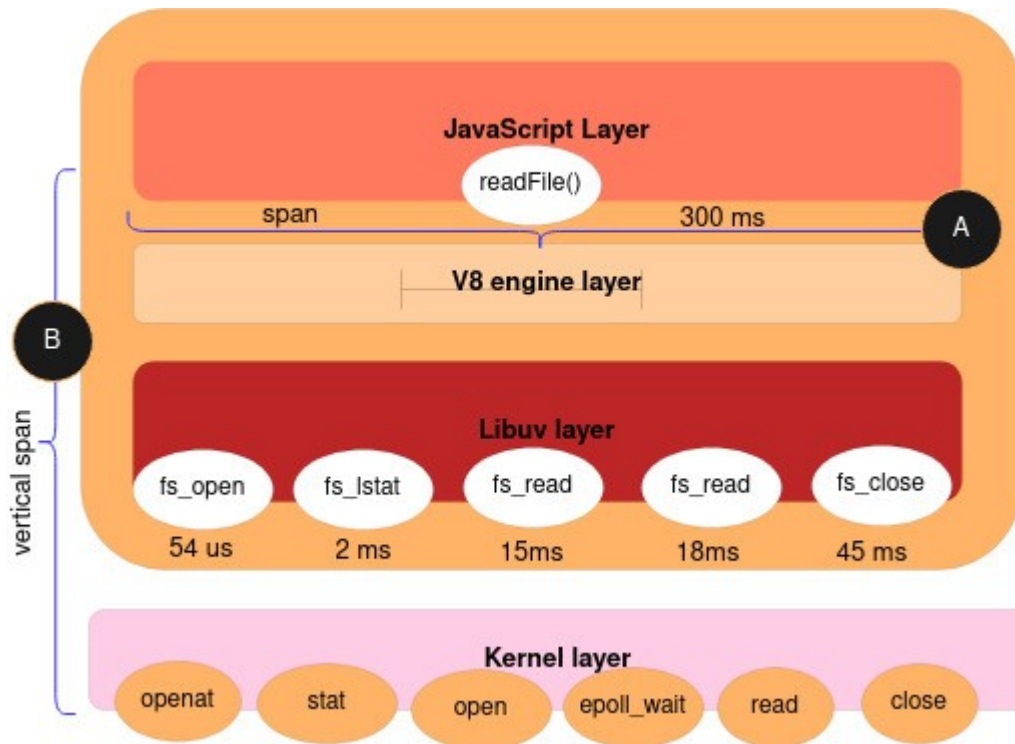


Introduction

NodeCompass approach

New paradigm:

- Introduce Vertical Span Model (VSM)

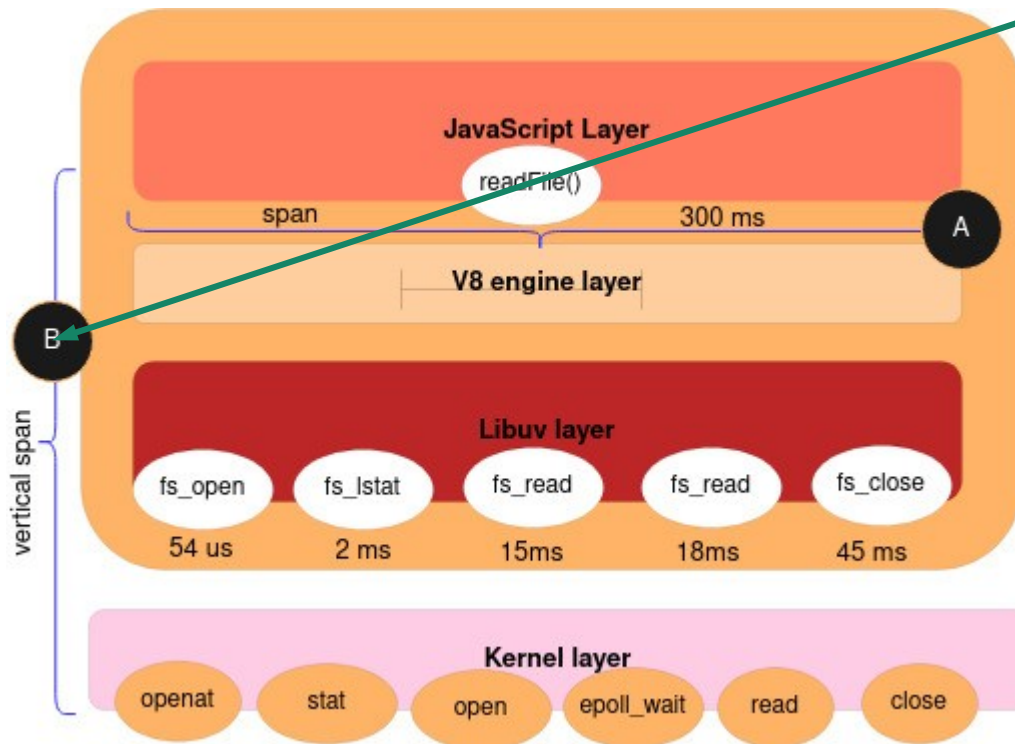


Introduction

NodeCompass approach

New paradigm:

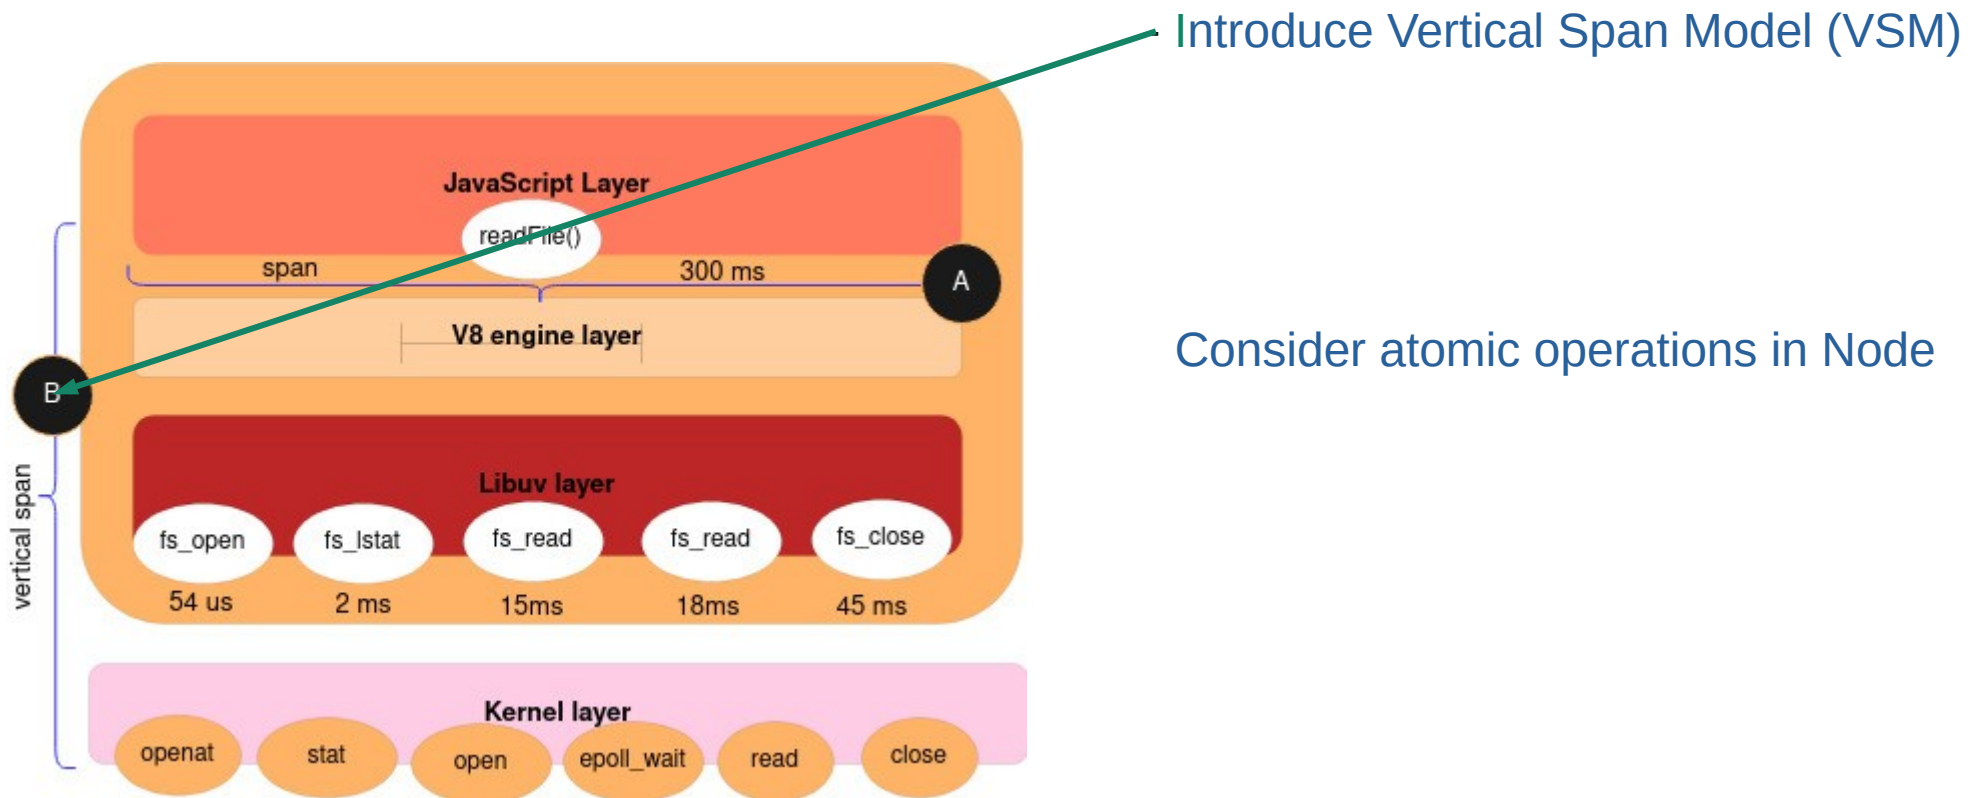
Introduce Vertical Span Model (VSM)



Introduction

NodeCompass approach

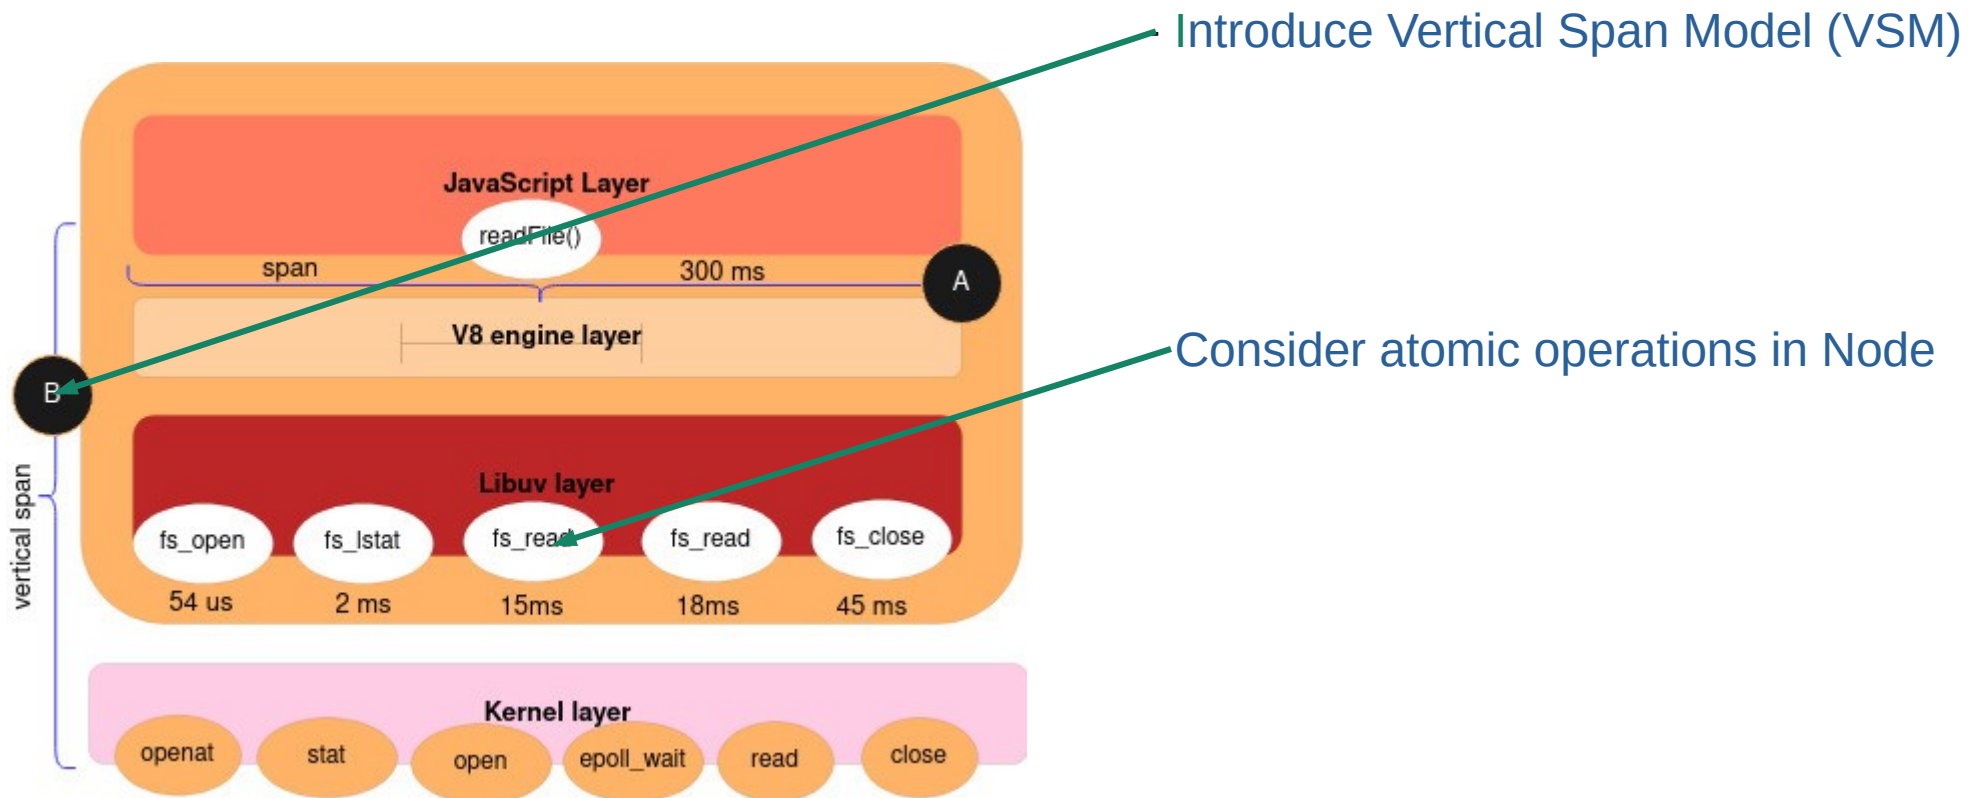
New paradigm:



Introduction

NodeCompass approach

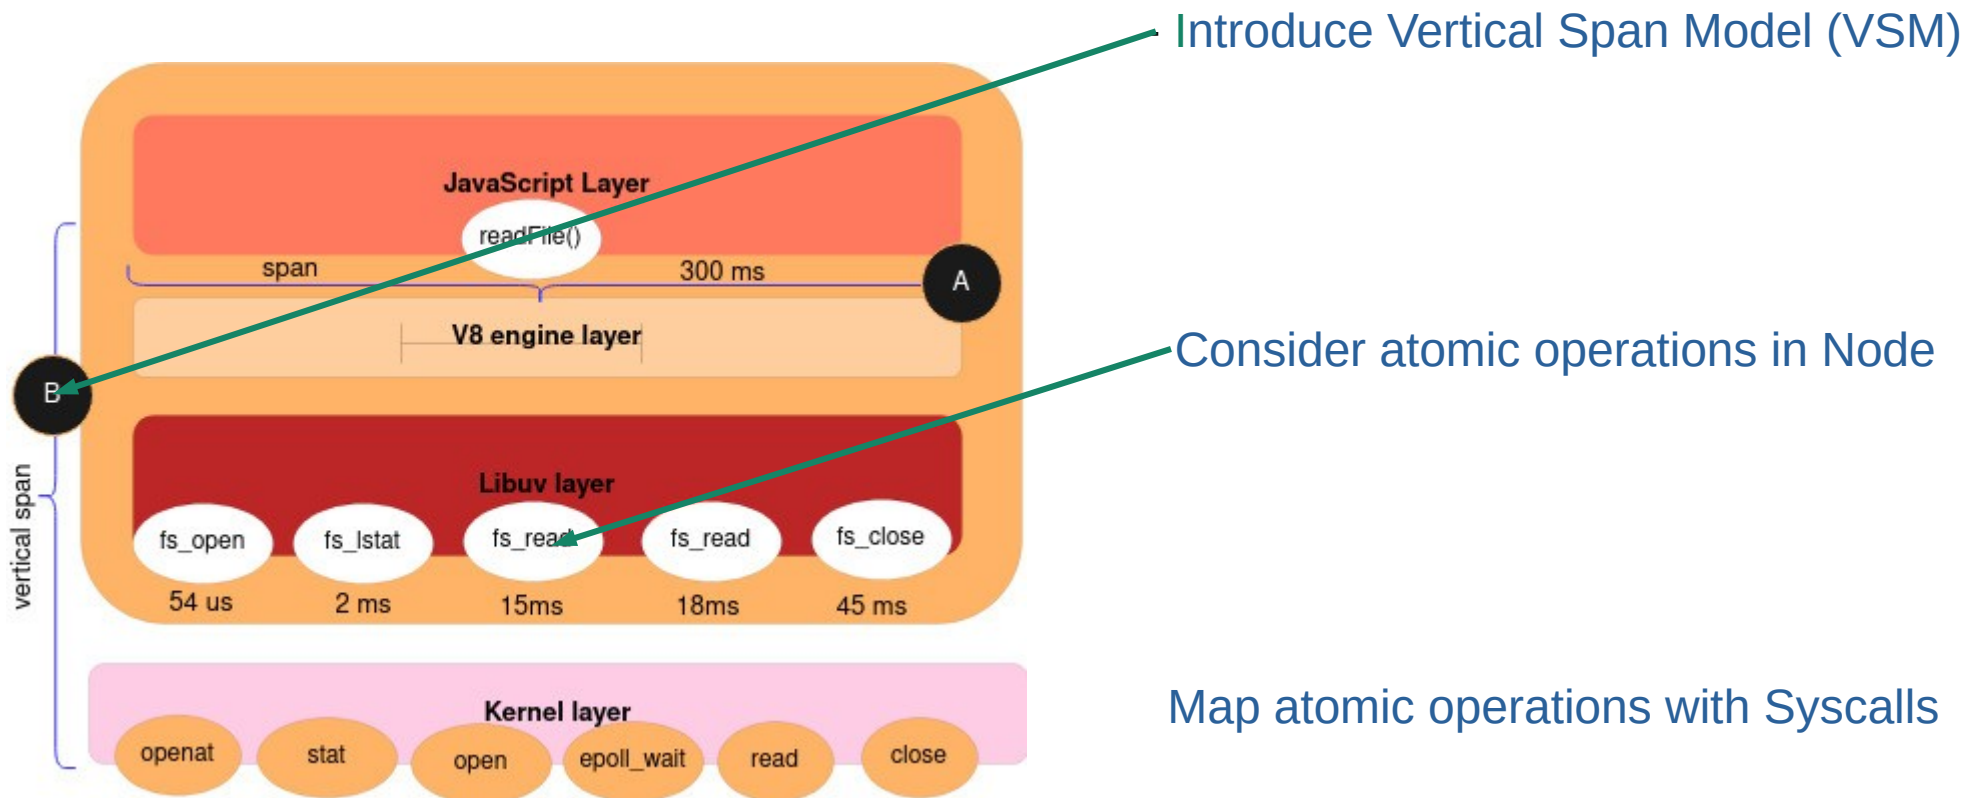
New paradigm:



Introduction

NodeCompass approach

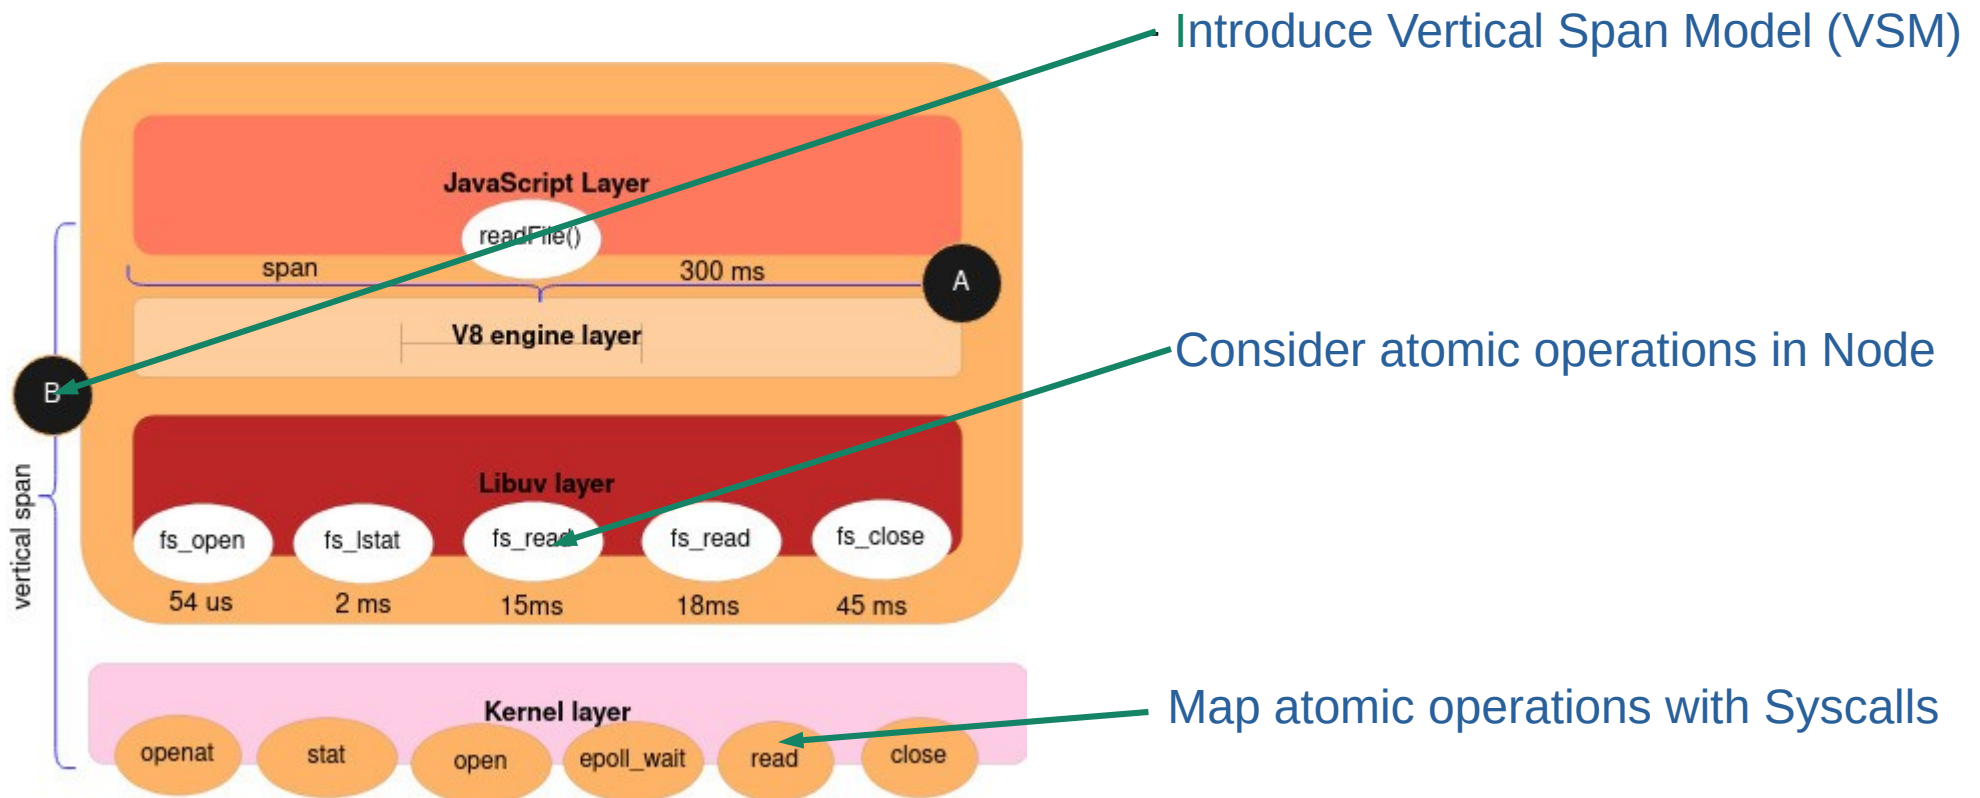
New paradigm:



Introduction

NodeCompass approach

New paradigm:



Introduction

The Vertical Span Model: The new paradigm

Introduction

The Vertical Span Model: The new paradigm

Represent a request with all its underlying atomic operations

Introduction

The Vertical Span Model: The new paradigm

Represent a request with all its underlying atomic operations

Expose the application flow and the request underlying paths

Introduction

The Vertical Span Model: The new paradigm

Represent a request with all its underlying atomic operations

Expose the application flow and the request underlying paths

Expose sequence of events in each layer with its latency

Introduction

The Vertical Span Model: The new paradigm

Represent a request with all its underlying atomic operations

Expose the application flow and the request underlying paths

Expose sequence of events in each layer with its latency

Performance analysis with high granularity degree

Introduction

The Vertical Span Model: The new paradigm

Represent a request with all its underlying atomic operations

Expose the application flow and the request underlying paths

Expose sequence of events in each layer with its latency

Performance analysis with high granularity degree

Pinpoint performance bottlenecks, system bugs, Root cause

Introduction

The Vertical Span Model: The new paradigm

Represent a request with all its underlying atomic operations

Expose the application flow and the request underlying paths

Expose sequence of events in each layer with its latency

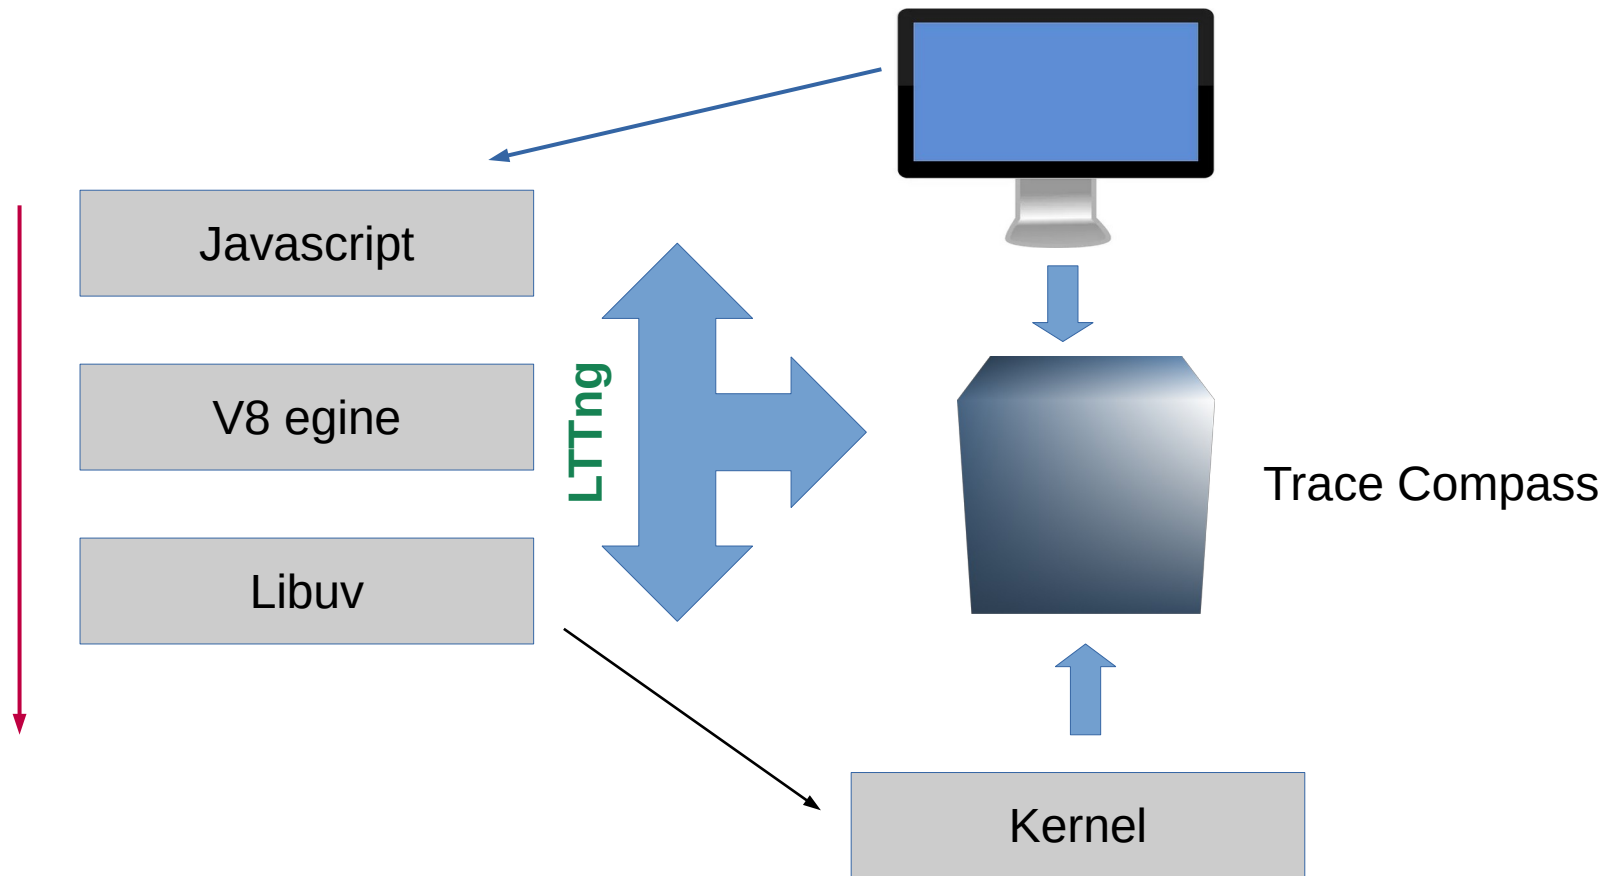
Performance analysis with high granularity degree

Pinpoint performance bottlenecks, system bugs, Root cause

Expose race Conditions in Node.js

Introduction

■ System Architecture



Architecture

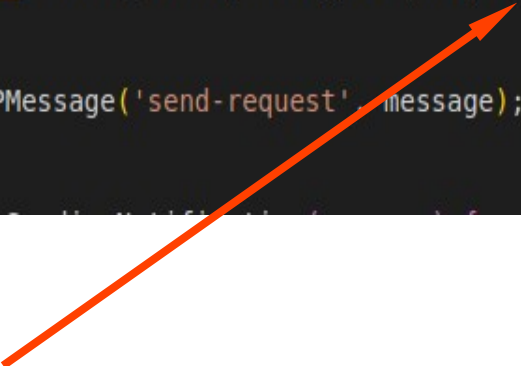
User level-Instrumentation Example

```
565 }
566 function traceSendingRequest(message) {
567     if (trace === Trace.Off || !tracer) {
568         return;
569     }
570     if (traceFormat === TraceFormat.Text) {
571         let data = undefined;
572         if (trace === Trace.Verbose && message.params) {
573             data = `Params: ${JSON.stringify(message.params, null, 4)}\n\n`;
574         }
575         //tracer.log(`Sending request '${message.method}' - (${message.id})'.`, data);
576         calculate.send_event(message.id, async_hooks.executionAsyncId(), "js_open_" + message.method + '$server');
577     }
578     else {
579         logLSPMessage('send-request', message);
580     }
581 }
```

Architecture

User level-Instrumentation Example

```
565 }
566 function traceSendingRequest(message) {
567     if (trace === Trace.Off || !tracer) {
568         return;
569     }
570     if (traceFormat === TraceFormat.Text) {
571         let data = undefined;
572         if (trace === Trace.Verbose && message.params) {
573             data = `Params: ${JSON.stringify(message.params, null, 4)}\n\n`;
574         }
575         //tracer.log(`Sending request '${message.method}' - (${message.id})'.`, data);
576         calculate.send_event(message.id, async_hooks.executionAsyncId(), "js_open_" + message.method + '$server');
577     }
578     else {
579         logLSPMessage('send-request', message);
580     }
581 }
```



Architecture

User level-Instrumentation Example

```
565 }
566 function traceSendingRequest(message) {
567     if (trace === Trace.Off || !tracer) {
568         return;
569     }
570     if (traceFormat === TraceFormat.Text) {
571         let data = undefined;
572         if (trace === Trace.Verbose && message.params) {
573             data = `Params: ${JSON.stringify(message.params, null, 4)}\n\n`;
574         }
575         //tracer.log(`Sending request '${message.method}' - (${message.id})'.`, data);
576         calculate.send_event(message.id, async_hooks.executionAsyncId(), "js_open_" + message.method + '$server');
577     }
578     else {
579         logLSPMessage('send-request', message);
580     }
581 }
```

Entry point (RPC)

Architecture

User level-Instrumentation Example

```
565 }
566 function traceSendingRequest(message) {
567     if (trace === Trace.Off || !tracer) {
568         return;
569     }
570     if (traceFormat === TraceFormat.Text) {
571         let data = undefined;
572         if (trace === Trace.Verbose && message.params) {
573             data = `Params: ${JSON.stringify(message.params, null, 4)}\n\n`;
574         }
575         //tracer.log(`Sending request '${message.method}' - (${message.id})',`, data);
576         calculate.send_event(message.id, async_hooks.executionAsyncId(), "js_open_" + message.method + '$server');
577     }
578     else {
579         logLSPMessage('send-request', message);
580     }
581 }
```

Entry point (RPC)

Architecture

User-Level Instrumentation Example

```
664 function traceReceivedResponse(message, responsePromise) {
665     if (trace === Trace.Off || !tracer) {
666         return;
667     }
668     if (traceFormat === TraceFormat.Text) {
669         let data = undefined;
670         if (trace === Trace.Verbose) {
671             if (message.error && message.error.data) {
672                 data = `Error data: ${JSON.stringify(message.error.data, null, 4)}\n\n`;
673             }
674             else {
675                 if (message.result) {
676                     data = `Result: ${JSON.stringify(message.result, null, 4)}\n\n`;
677                 }
678                 else if (message.error === void 0) {
679                     data = `No result returned.\n\n`;
680                 }
681             }
682         }
683         if (responsePromise) {
684             let error = message.error ? ` Request failed: ${message.error.message} (${message.error.code}).` : '';
685             //tracer.log(`Received response '${responsePromise.method}' - (${message.id})' in ${Date.now() - responseProm
686             calculate.send_event(message.id, async_hooks.executionAsyncId(), 'js_exit$server');
687         }
688     }
```


Architecture

User-Level Instrumentation Example

```
664 function traceReceivedResponse(message, responsePromise) {
665     if (trace === Trace.Off || !tracer) {
666         return;
667     }
668     if (traceFormat === TraceFormat.Text) {
669         let data = undefined;
670         if (trace === Trace.Verbose) {
671             if (message.error && message.error.data) {
672                 data = `Error data: ${JSON.stringify(message.error.data, null, 4)}\n\n`;
673             }
674             else {
675                 if (message.result) {
676                     data = `Result: ${JSON.stringify(message.result, null, 4)}\n\n`;
677                 }
678                 else if (message.error === void 0) {
679                     data = `No result returned.\n\n`;
680                 }
681             }
682         }
683         if (responsePromise) {
684             let error = message.error ? `Request failed: ${message.error.message} (${message.error.code}).` : '';
685             //tracer.log(`Received response '${responsePromise.method}' - (${message.id})' in ${Date.now() - responseProm
686             calculate.send_event(message.id, async_hooks.executionAsyncId(), 'js_exit$server');
687         }
688     }
```

Exit Point (RPC)

Architecture

User-Level Instrumentation Example

```
664 function traceReceivedResponse(message, responsePromise) {
665     if (trace === Trace.Off || !tracer) {
666         return;
667     }
668     if (traceFormat === TraceFormat.Text) {
669         let data = undefined;
670         if (trace === Trace.Verbose) {
671             if (message.error && message.error.data) {
672                 data = `Error data: ${JSON.stringify(message.error.data, null, 4)}\n\n`;
673             }
674             else {
675                 if (message.result) {
676                     data = `Result: ${JSON.stringify(message.result, null, 4)}\n\n`;
677                 }
678                 else if (message.error === void 0) {
679                     data = `No result returned.\n\n`;
680                 }
681             }
682         }
683         if (responsePromise) {
684             let error = message.error ? `Request failed: ${message.error.message} (${message.error.code}).` : '';
685             //tracer.log(`Received response '${responsePromise.method}' - (${message.id})' in ${Date.now() - responseProm`
686             calculate.send_event(message.id, async_hooks.executionAsyncId(), 'js_exit$server');
687         }
688     }
```

Exit Point (RPC)

Architecture

User-Level Instrumentation Example

```
664 function traceReceivedResponse(message, responsePromise) {
665   if (trace === Trace.Off || !tracer) {
666     return;
667   }
668   if (traceFormat === TraceFormat.Text) {
669     let data = undefined;
670     if (trace === Trace.Verbose) {
671       if (message.error && message.error.data) {
672         data = `Error data: ${JSON.stringify(message.error.data, null, 4)}\n\n`;
673       }
674       else {
675         if (message.result) {
676           data = `Result: ${JSON.stringify(message.result, null, 4)}\n\n`;
677         }
678         else if (message.error === void 0) {
679           data = 'No result returned.\n\n';
680         }
681       }
682     }
683     if (responsePromise) {
684       let error = message.error ? `Request failed: ${message.error.message} (${message.error.code}).` : '';
685       //tracer.log('Received response `${responsePromise.method}` - (${message.id})' in ${Date.now() - responseProm
686       calculate.send_event(message.id, async_hooks.executionAsyncId(), 'js_exit$server');
687     }
688   }
```

Exit Point (RPC)

Use case 1

Pinpointing process.NextTick Performance Issues

Use case 1

Pinpointing process.NexTick Performance Issues

Postpone function execution to the Event-loop next tick

Use case 1

Pinpointing process.NexTick Performance Issues

Postpone function execution to the Event-loop next tick

Queue must be exhausted before the Event-loop continues to next phases

Use case 1

Pinpointing process.NextTick Performance Issues

Postpone function execution to the Event-loop next tick

Queue must be exhausted before the Event-loop continues to next phases

Must be used with caution since it can slow or block the Event-loop

Nextick Performance degradation scenario

Use case 1

Pinpointing process.NextTick Performance Issues

Postpone function execution to the Event-loop next tick

Queue must be exhausted before the Event-loop continues to next phases

Must be used with caution since it can slow or block the Event-loop

Nextick Performance degradation scenario

An express server contacted on the rout /nexttick

Use case 1

Pinpointing process.NextTick Performance Issues

Postpone function execution to the Event-loop next tick

Queue must be exhausted before the Event-loop continues to next phases

Must be used with caution since it can slow or block the Event-loop

Nextick Performance degradation scenario

An express server contacted on the rout /nexttick

A client **A** contacting continuously the server at the route /nexttick

Use case 1

Pinpointing process.NextTick Performance Issues

Postpone function execution to the Event-loop next tick

Queue must be exhausted before the Event-loop continues to next phases

Must be used with caution since it can slow or block the Event-loop

Nextick Performance degradation scenario

An express server contacted on the rout /nexttick

A client **A** contacting continuously the server at the route /nexttick

A second client **B** contacting continuously the server at the route /checkState

Use case 1

Pinpointing process.NextTick Performance Issues

Postpone function execution to the Event-loop next tick

Queue must be exhausted before the Event-loop continues to next phases

Must be used with caution since it can slow or block the Event-loop

Nextick Performance degradation scenario

An express server contacted on the rout /nexttick

A client **A** contacting continuously the server at the route /nexttick

A second client **B** contacting continuously the server at the route /checkState

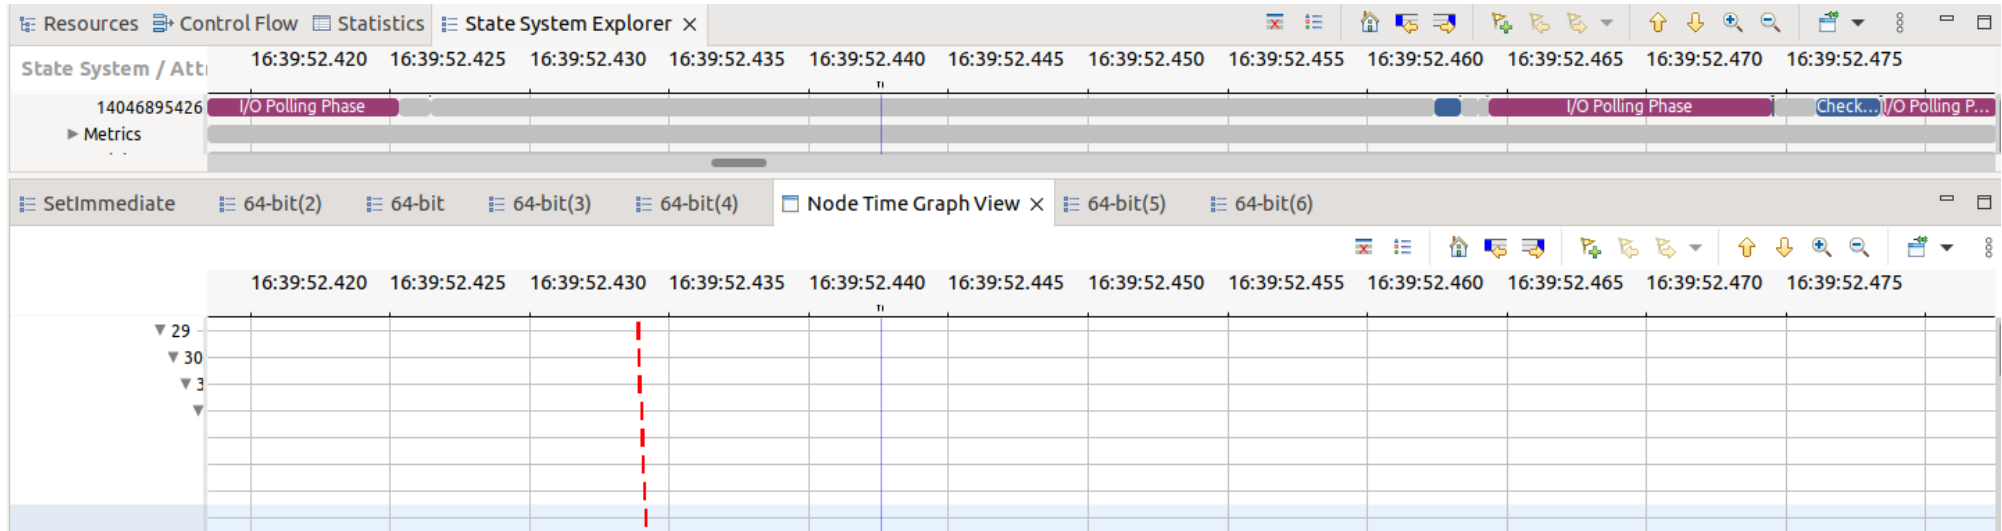
After a few seconds the server stops responding to B. A performance issue appears

Use case 1

Resulting flow:

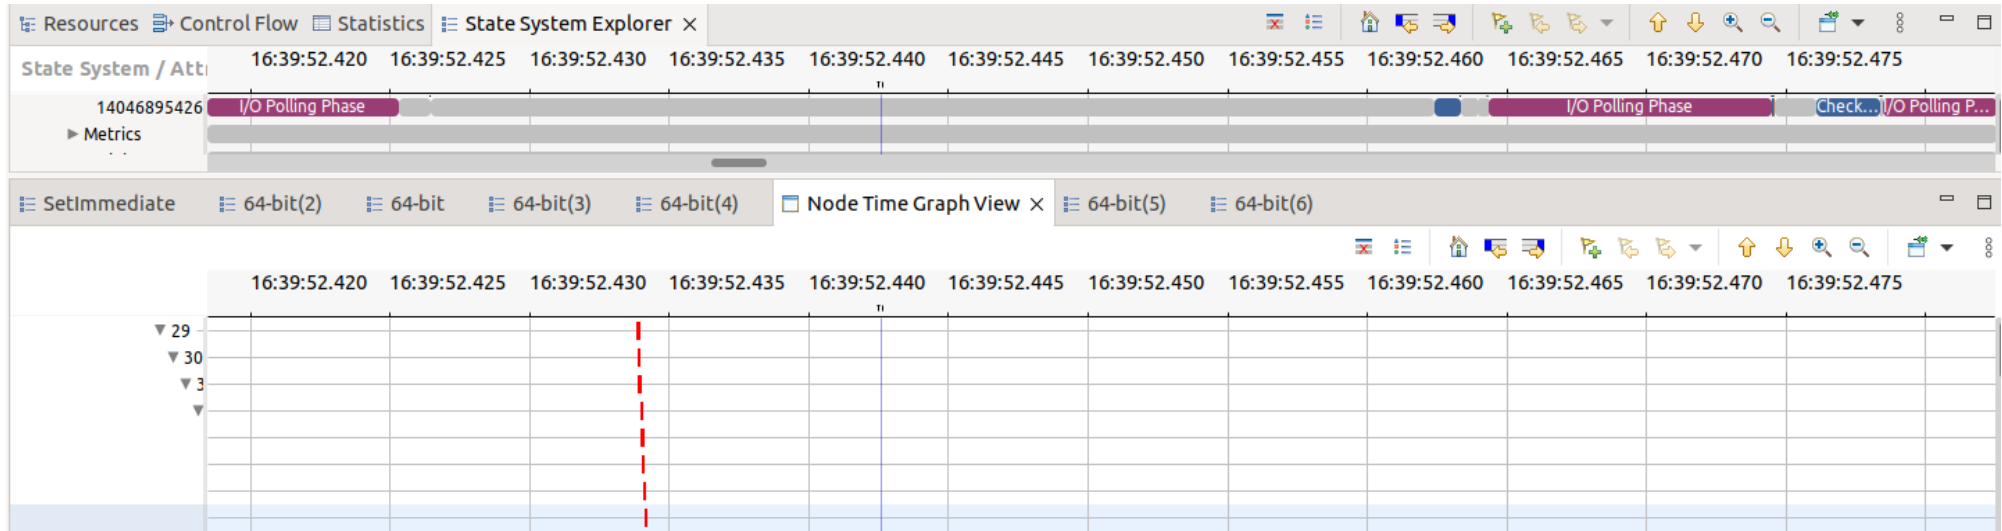
Use case 1

Resulting flow:



Use case 1

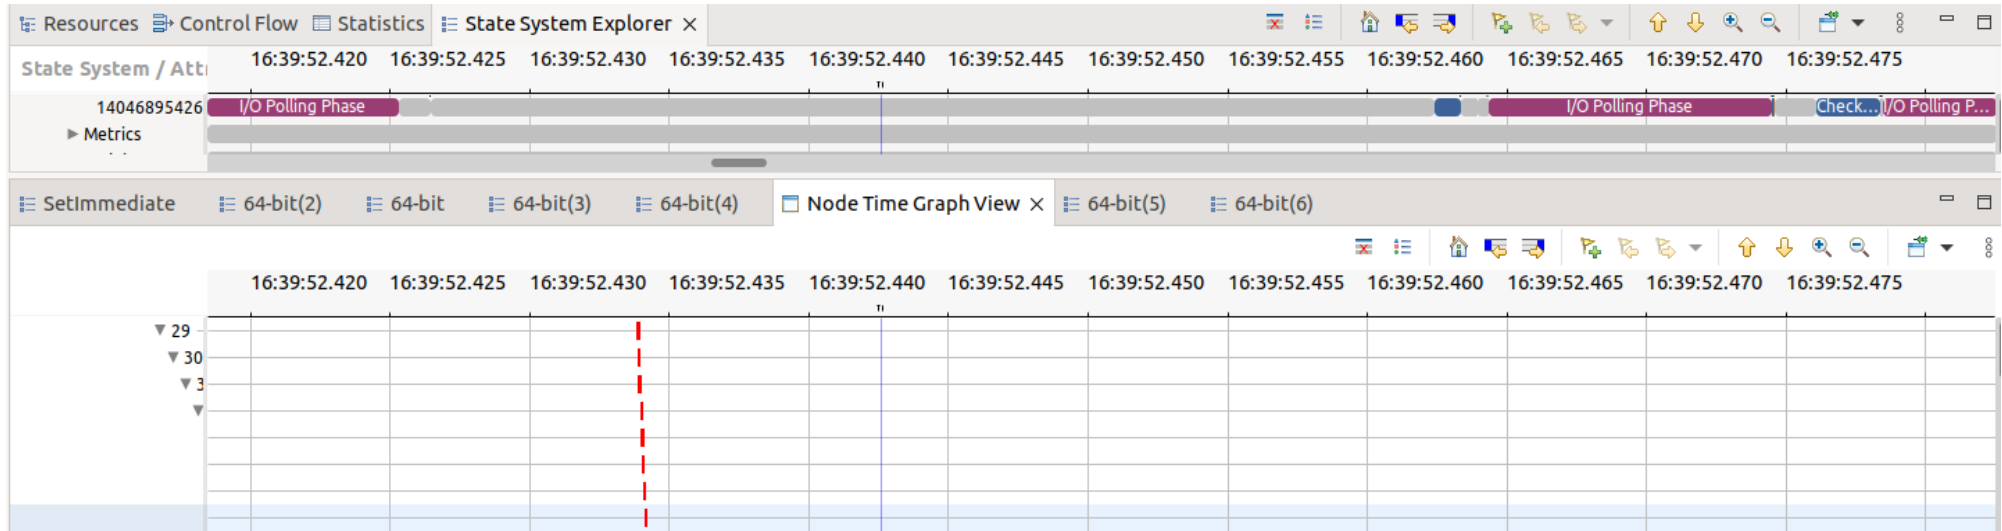
Resulting flow:



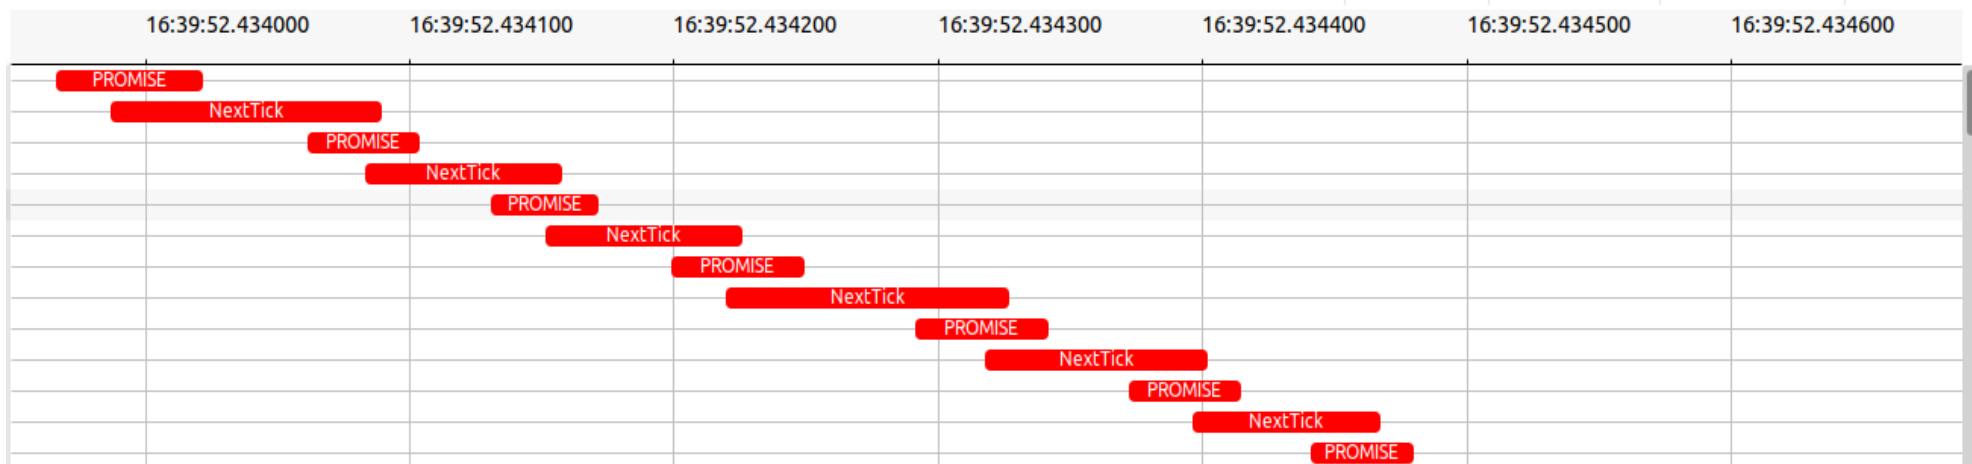
Let zoom on the view

Use case 1

Resulting flow:



Let zoom on the view

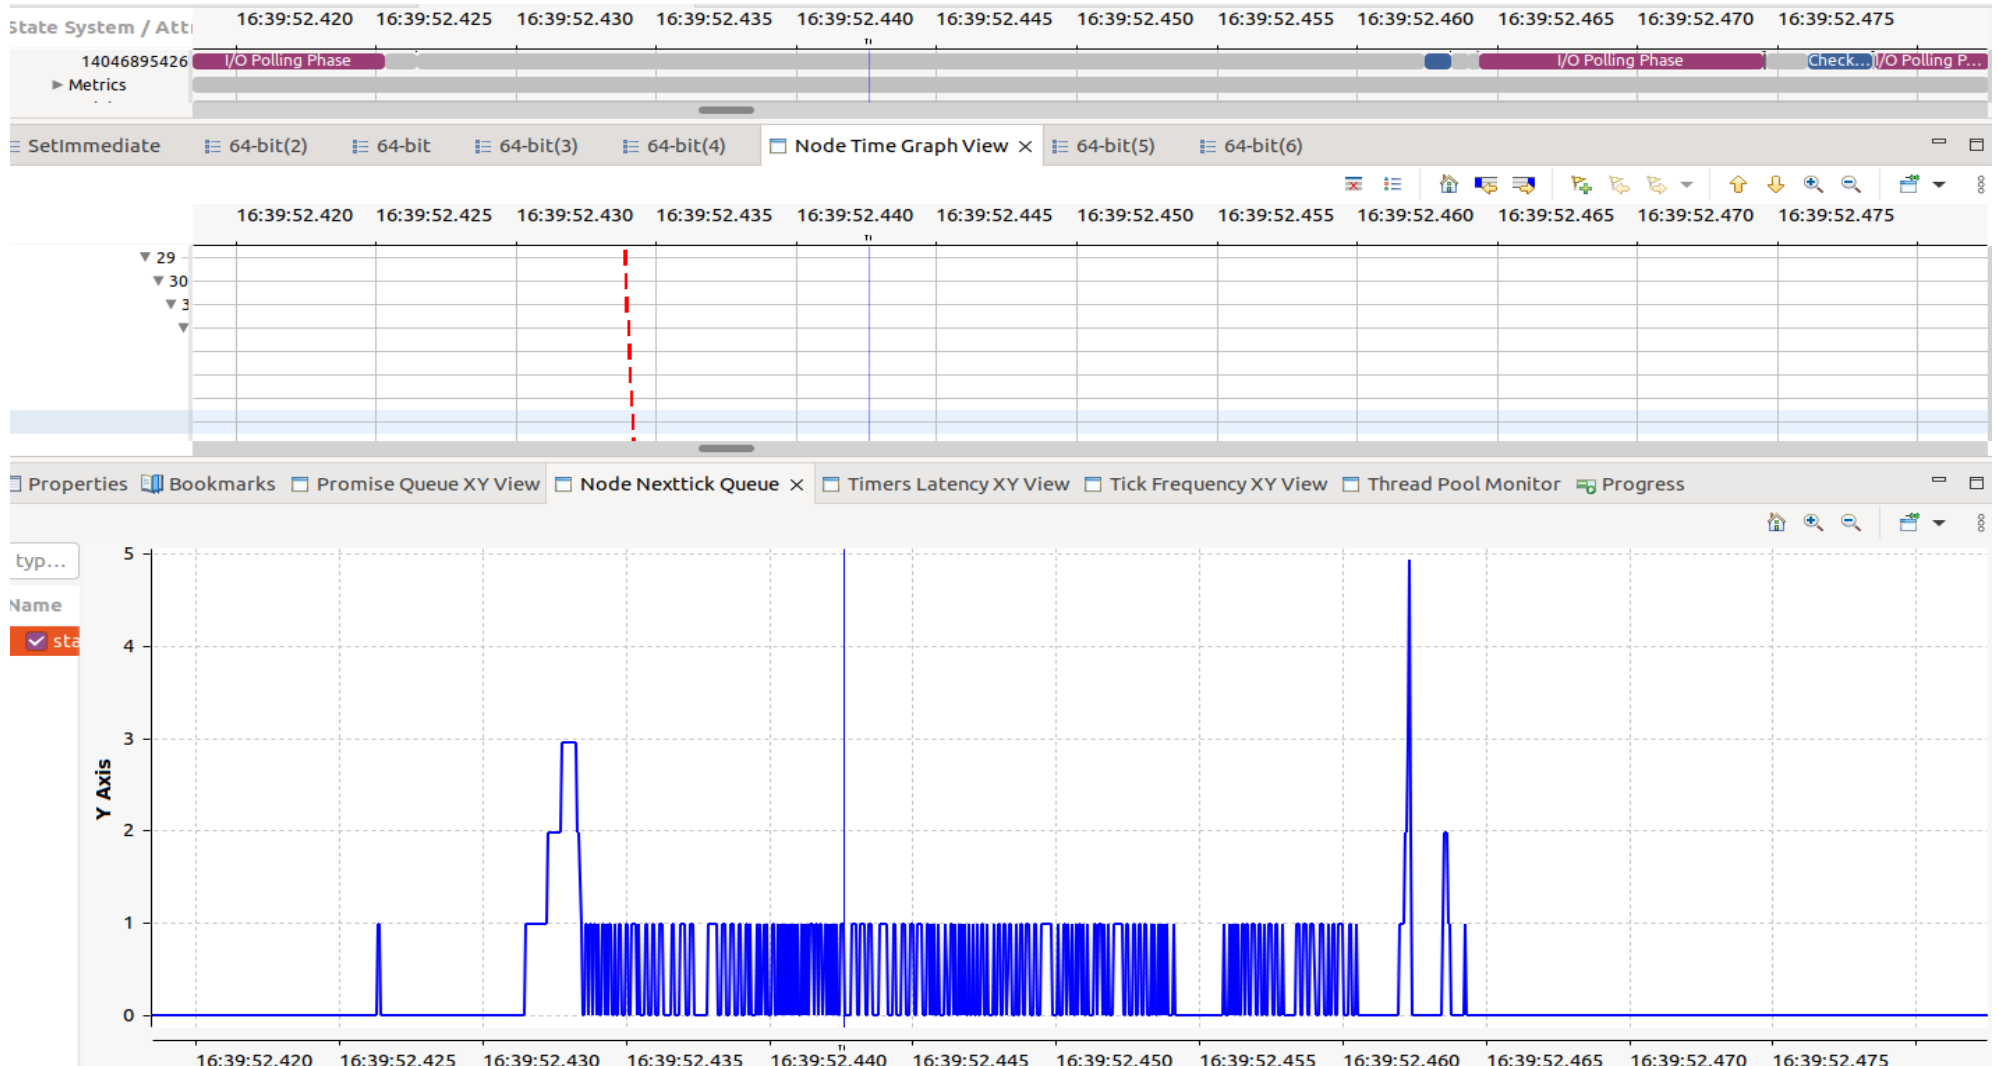


Use case 1

Results Interpretation

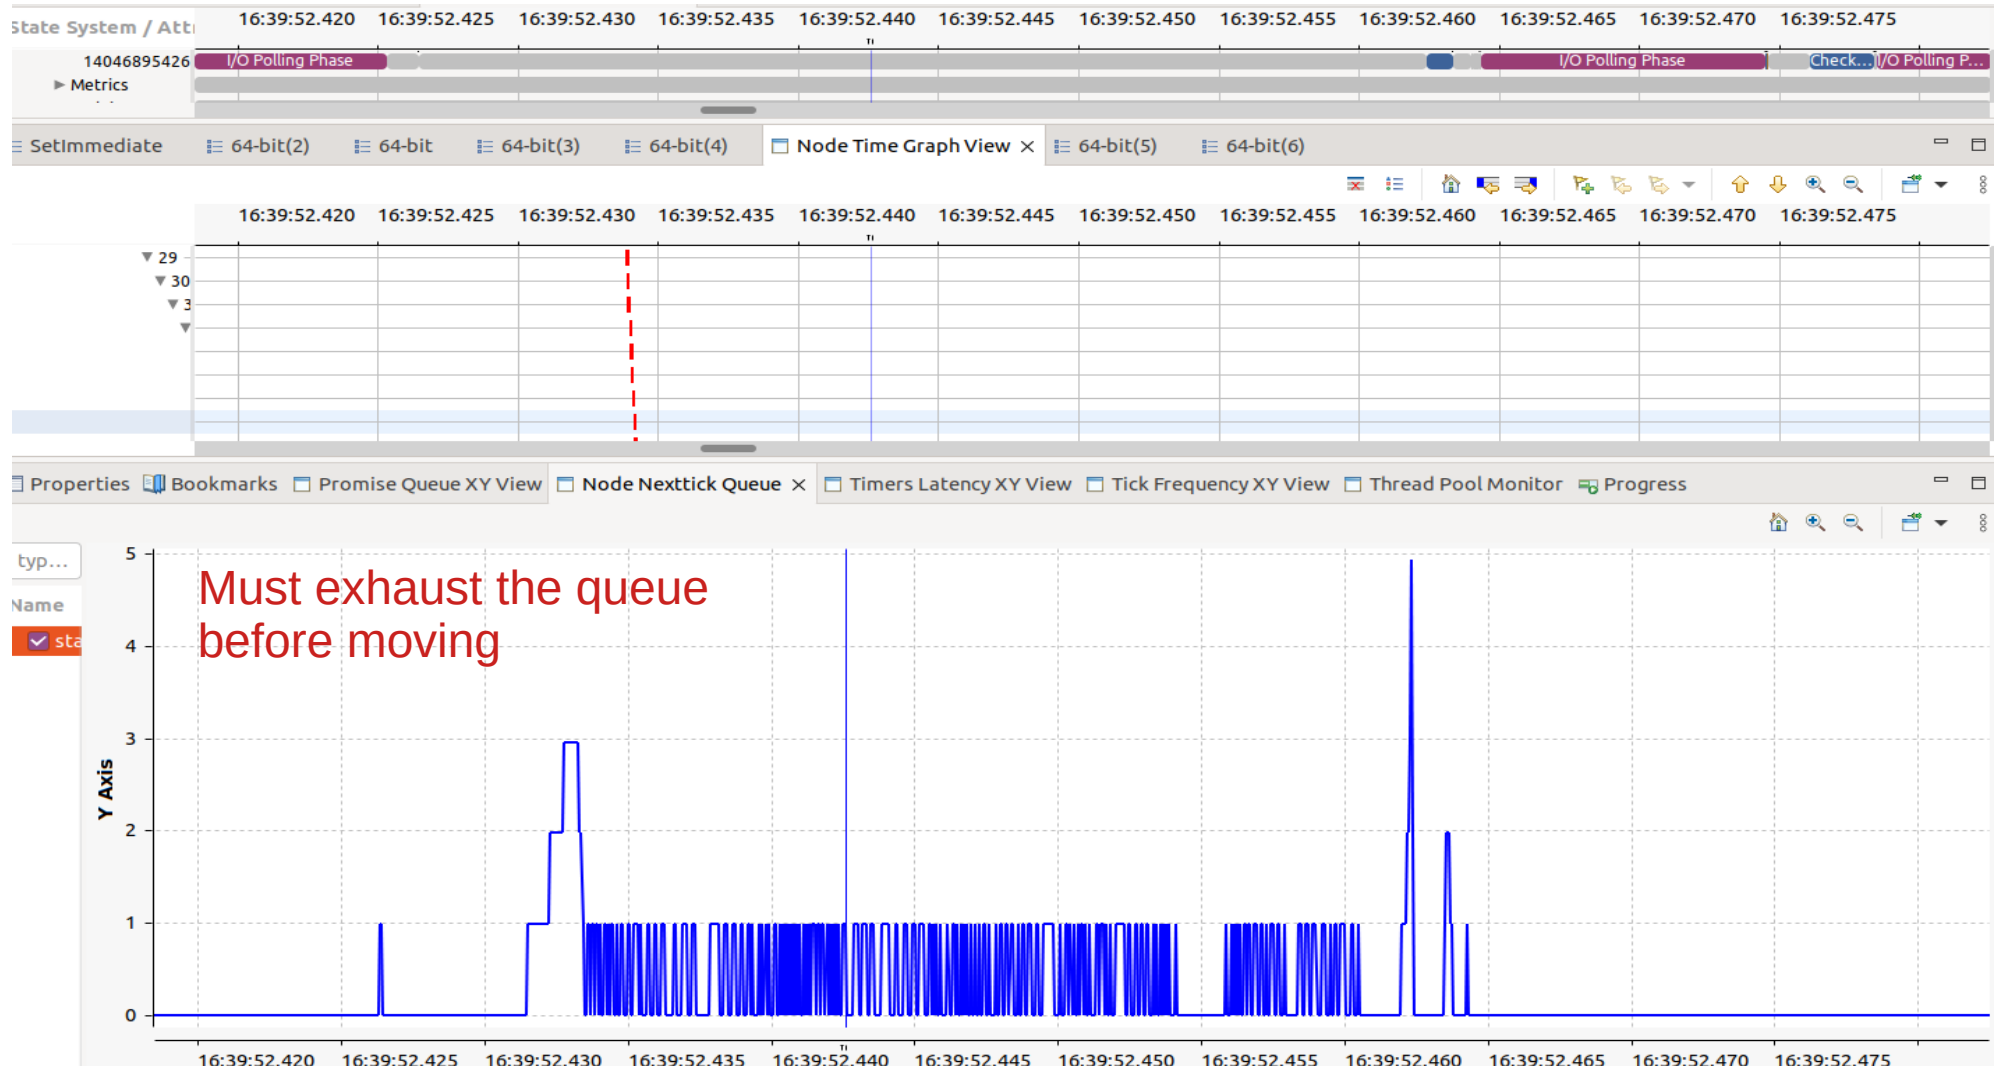
Use case 1

Results Interpretation



Use case 1

Results Interpretation



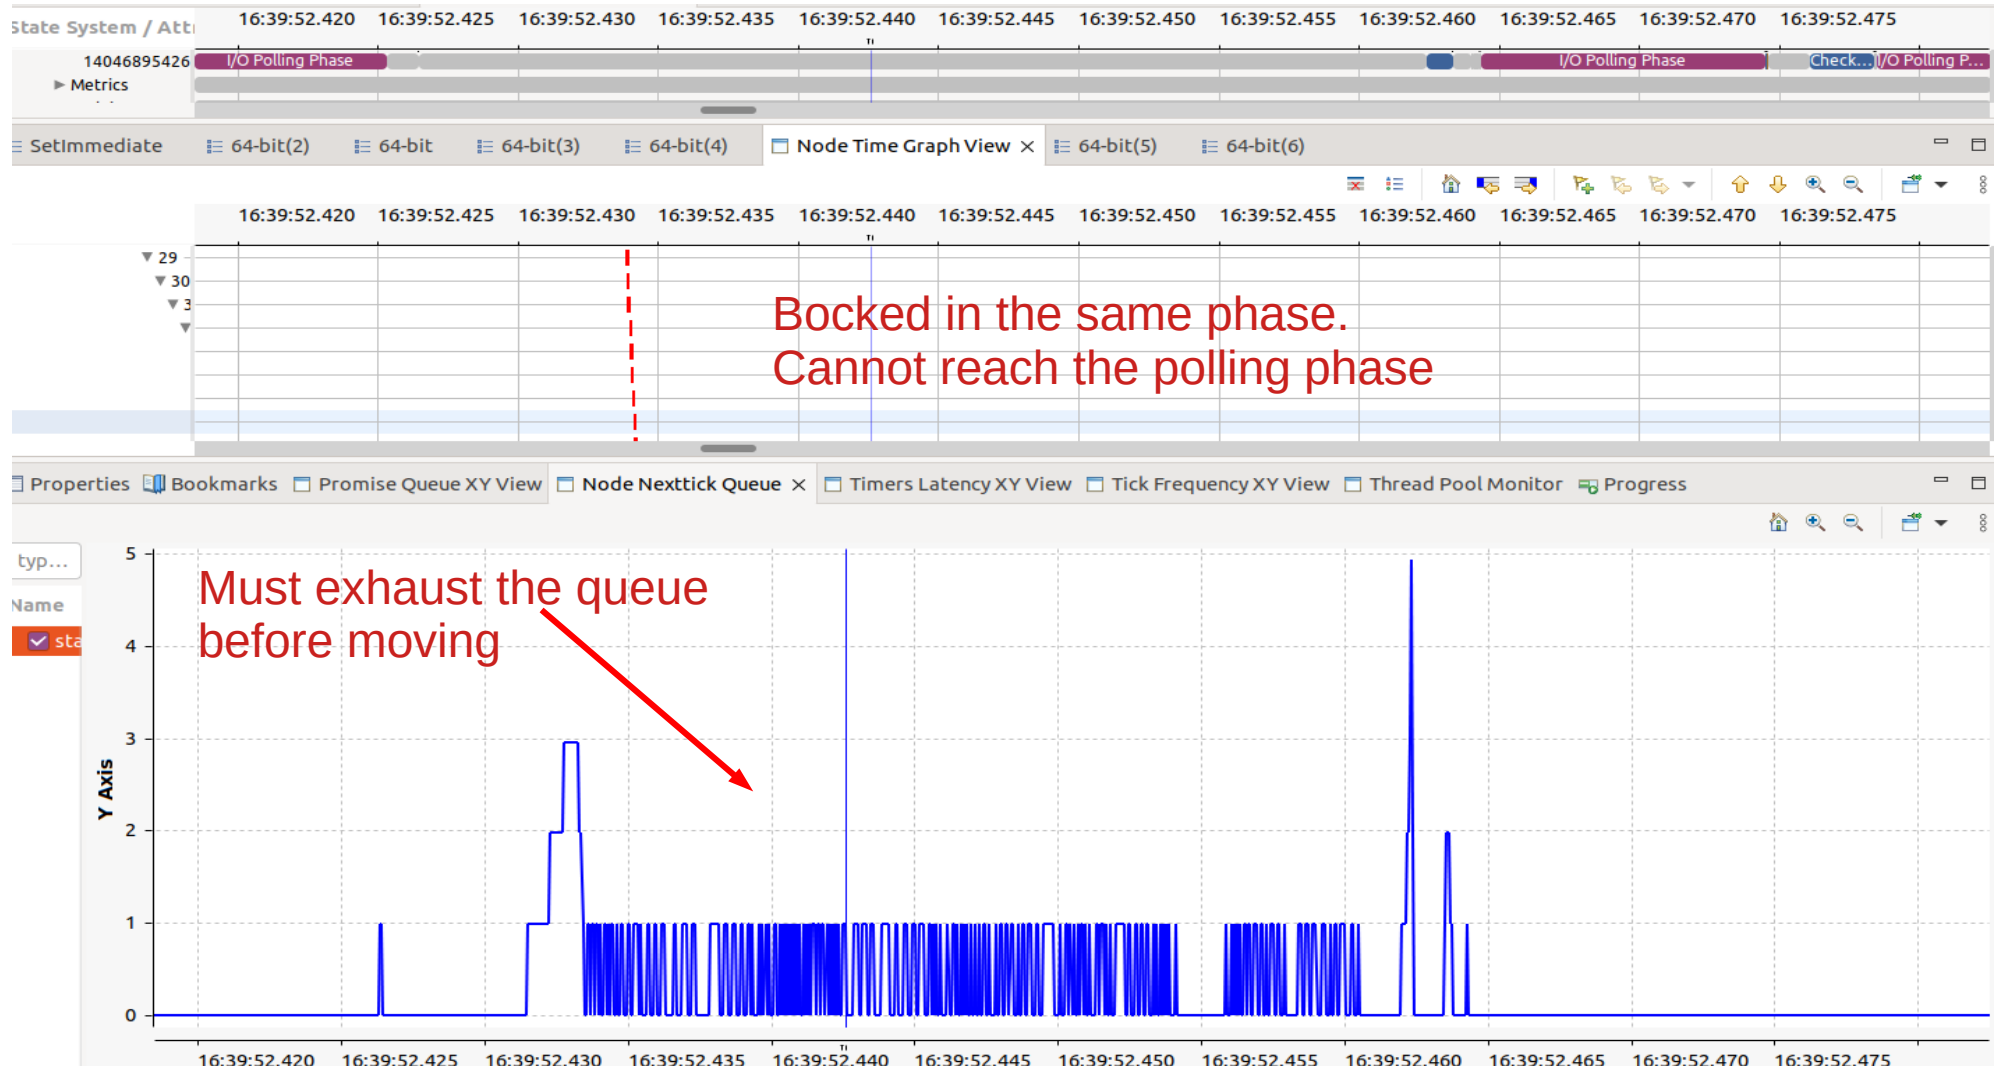
Use case 1

Results Interpretation



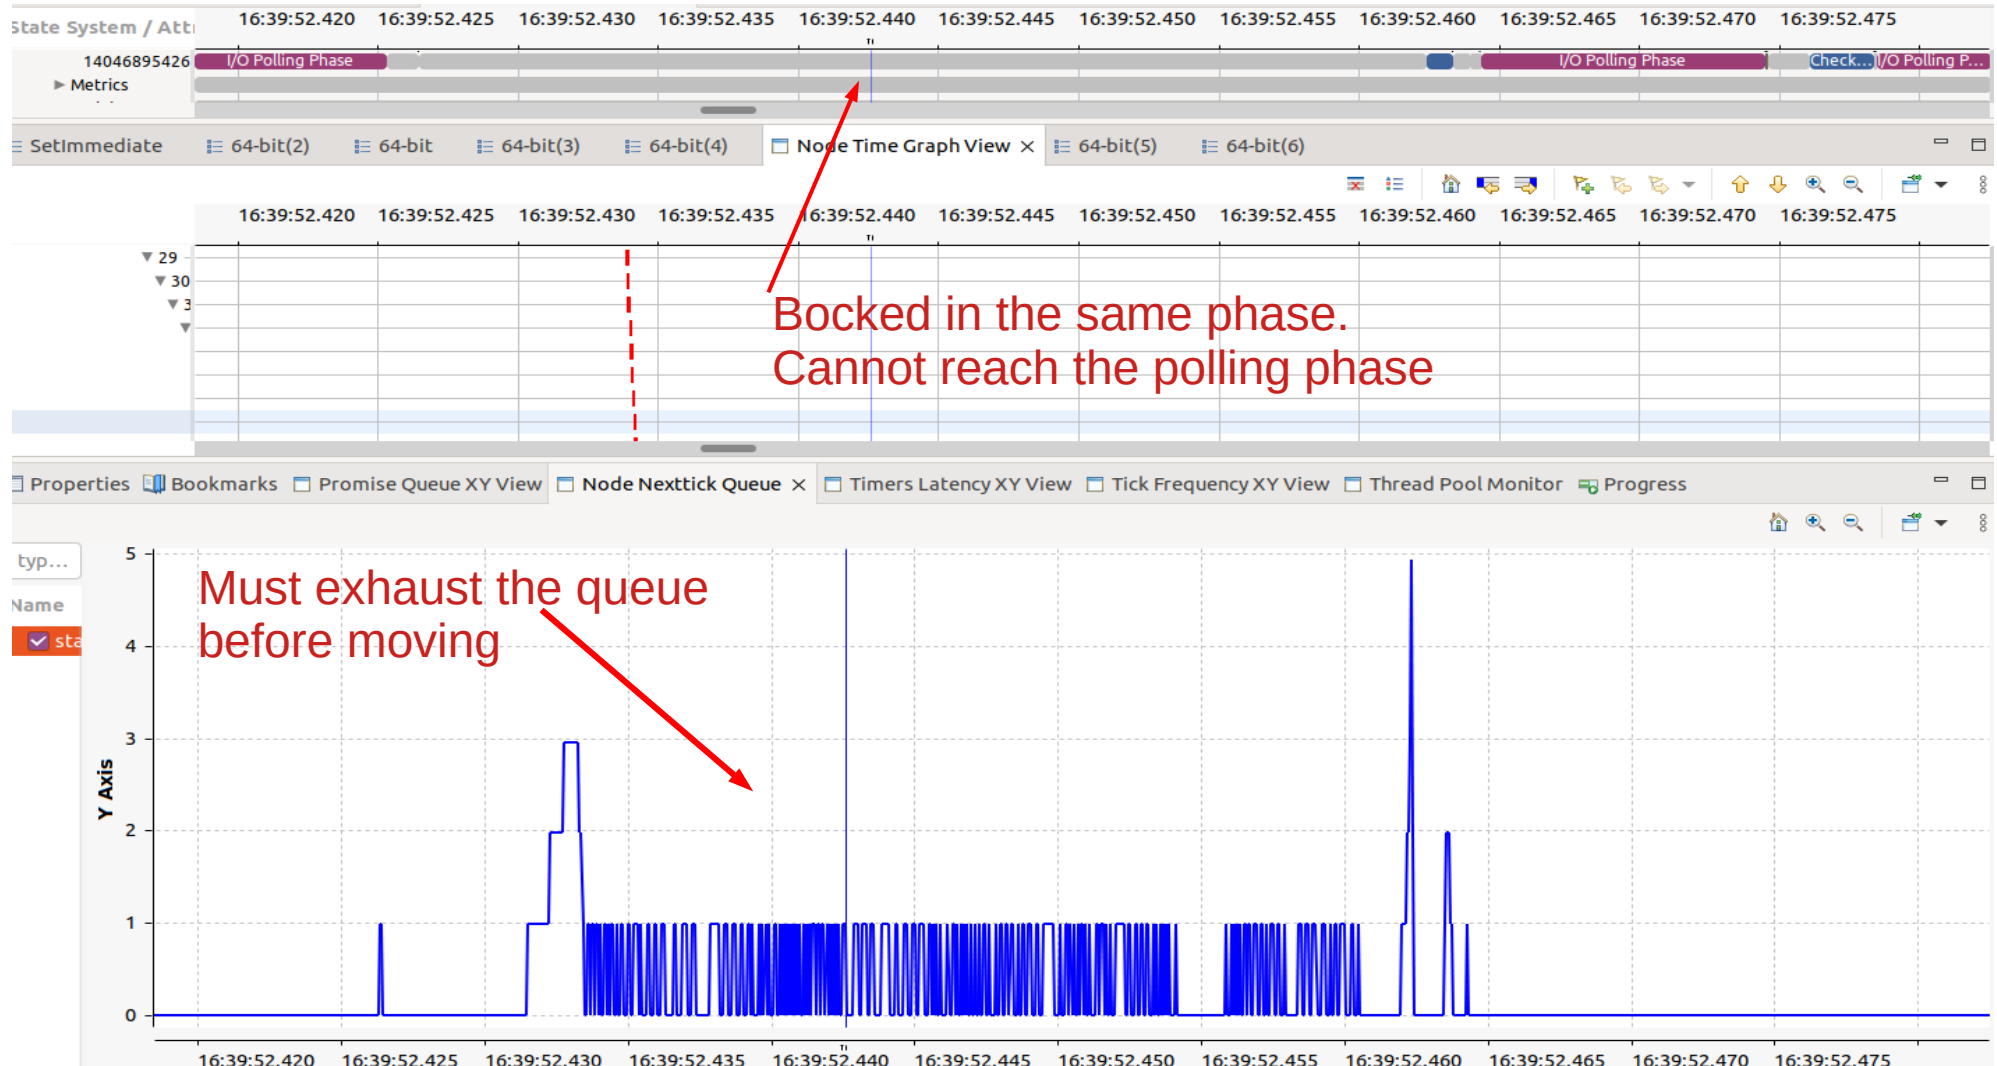
Use case 1

Results Interpretation



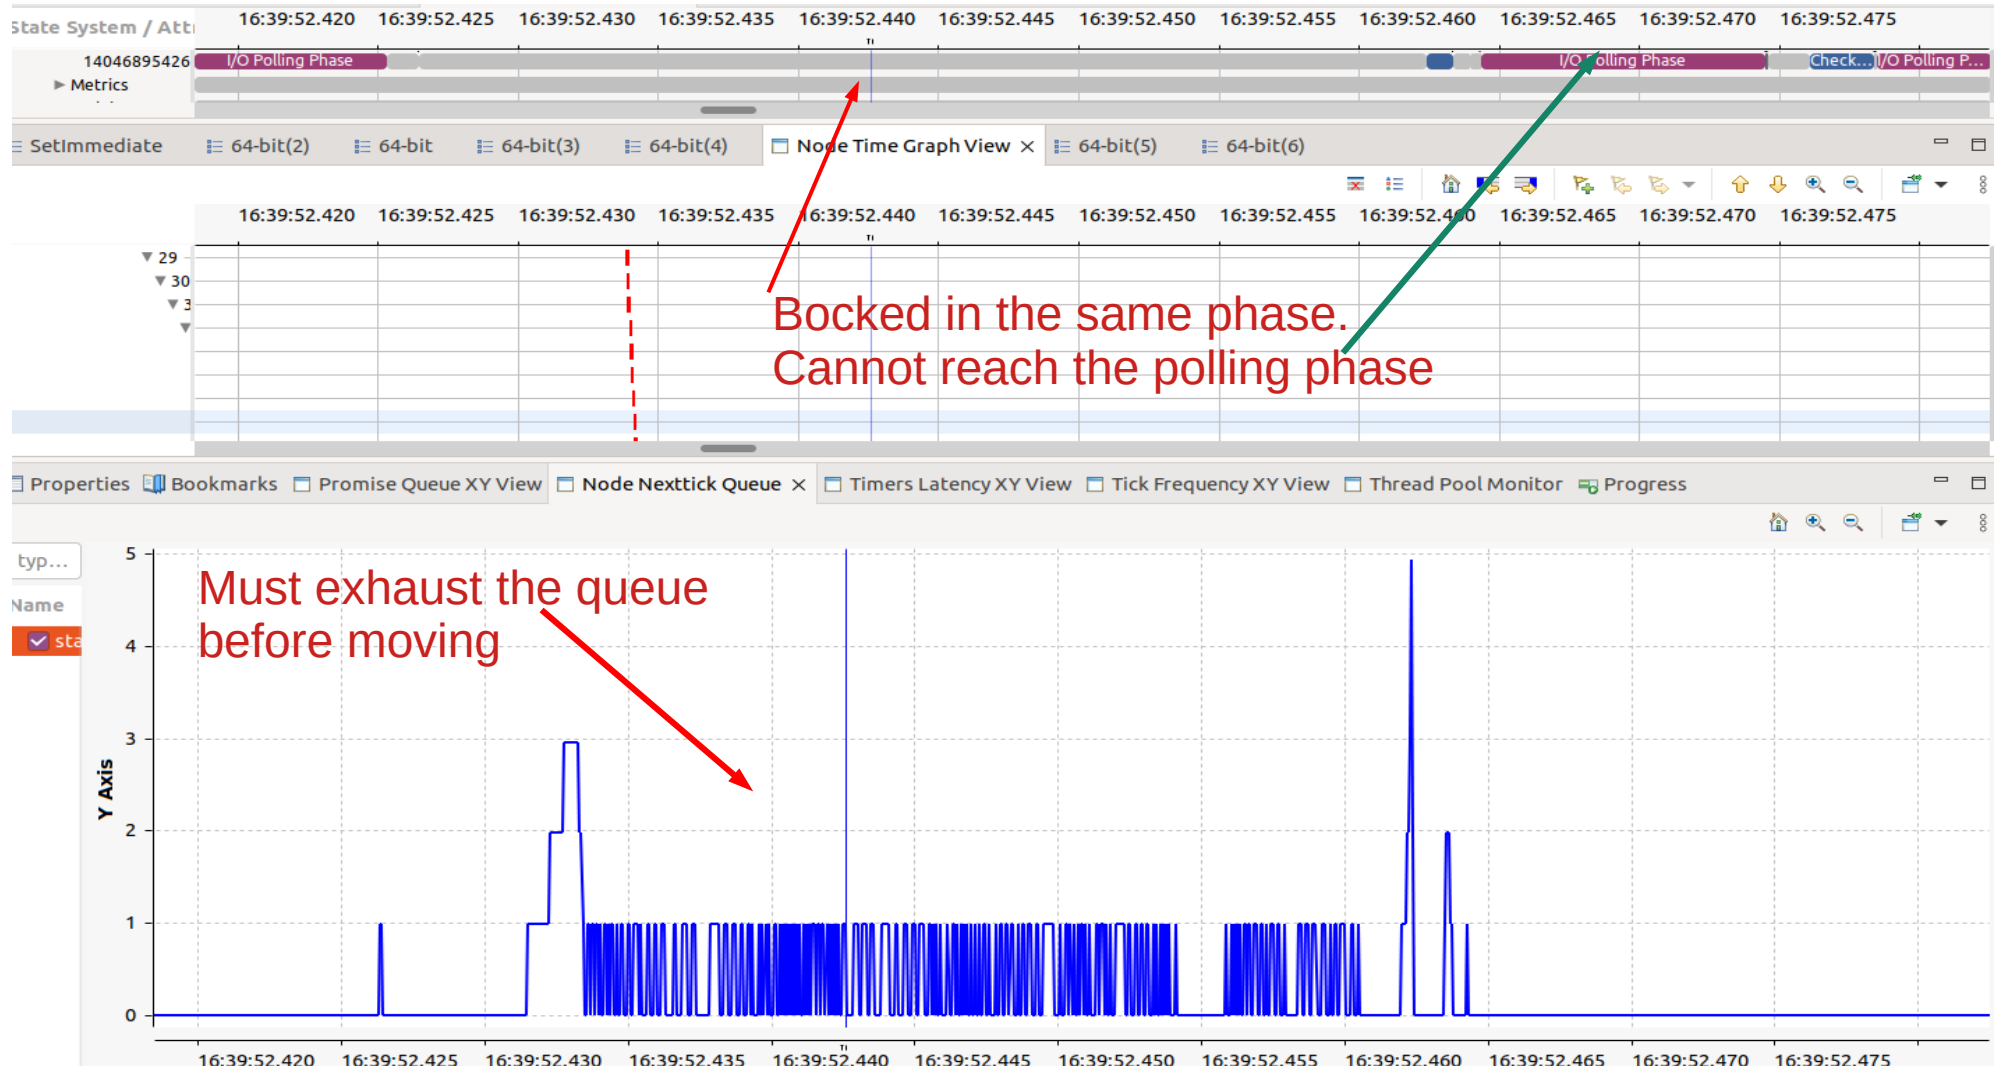
Use case 1

Results Interpretation



Use case 1

Results Interpretation



Use case 2

Asynchronous resources destruction bug

Use case 2

Asynchronous resources destruction bug

Each asynchronous resource has a lifetime

Use case 2

Asynchronous resources destruction bug

Each asynchronous resource has a lifetime

Internal mechanisms in Node.js track async. resources and destroy them after used

Use case 2

Asynchronous resources destruction bug

Each asynchronous resource has a lifetime

Internal mechanisms in Node.js track async. resources and destroy them after used

“Destroy” callback is invoked when cleaning the system from the resource

Test scenario

Use case 2

Asynchronous resources destruction bug

Each asynchronous resource has a lifetime

Internal mechanisms in Node.js track async. resources and destroy them after used

“Destroy” callback is invoked when cleaning the system from the resource

Test scenario

An express server contacted at the route /test

Use case 2

Asynchronous resources destruction bug

Each asynchronous resource has a lifetime

Internal mechanisms in Node.js track async. resources and destroy them after used

“Destroy” callback is invoked when cleaning the system from the resource

Test scenario

An express server contacted at the route /test

A client **A** contacting continuously the server on the route /test

Use case 2

Asynchronous resources destruction bug

Each asynchronous resource has a lifetime

Internal mechanisms in Node.js track async. resources and destroy them after used

“Destroy” callback is invoked when cleaning the system from the resource

Test scenario

An express server contacted at the route /test

A client **A** contacting continuously the server on the route /test

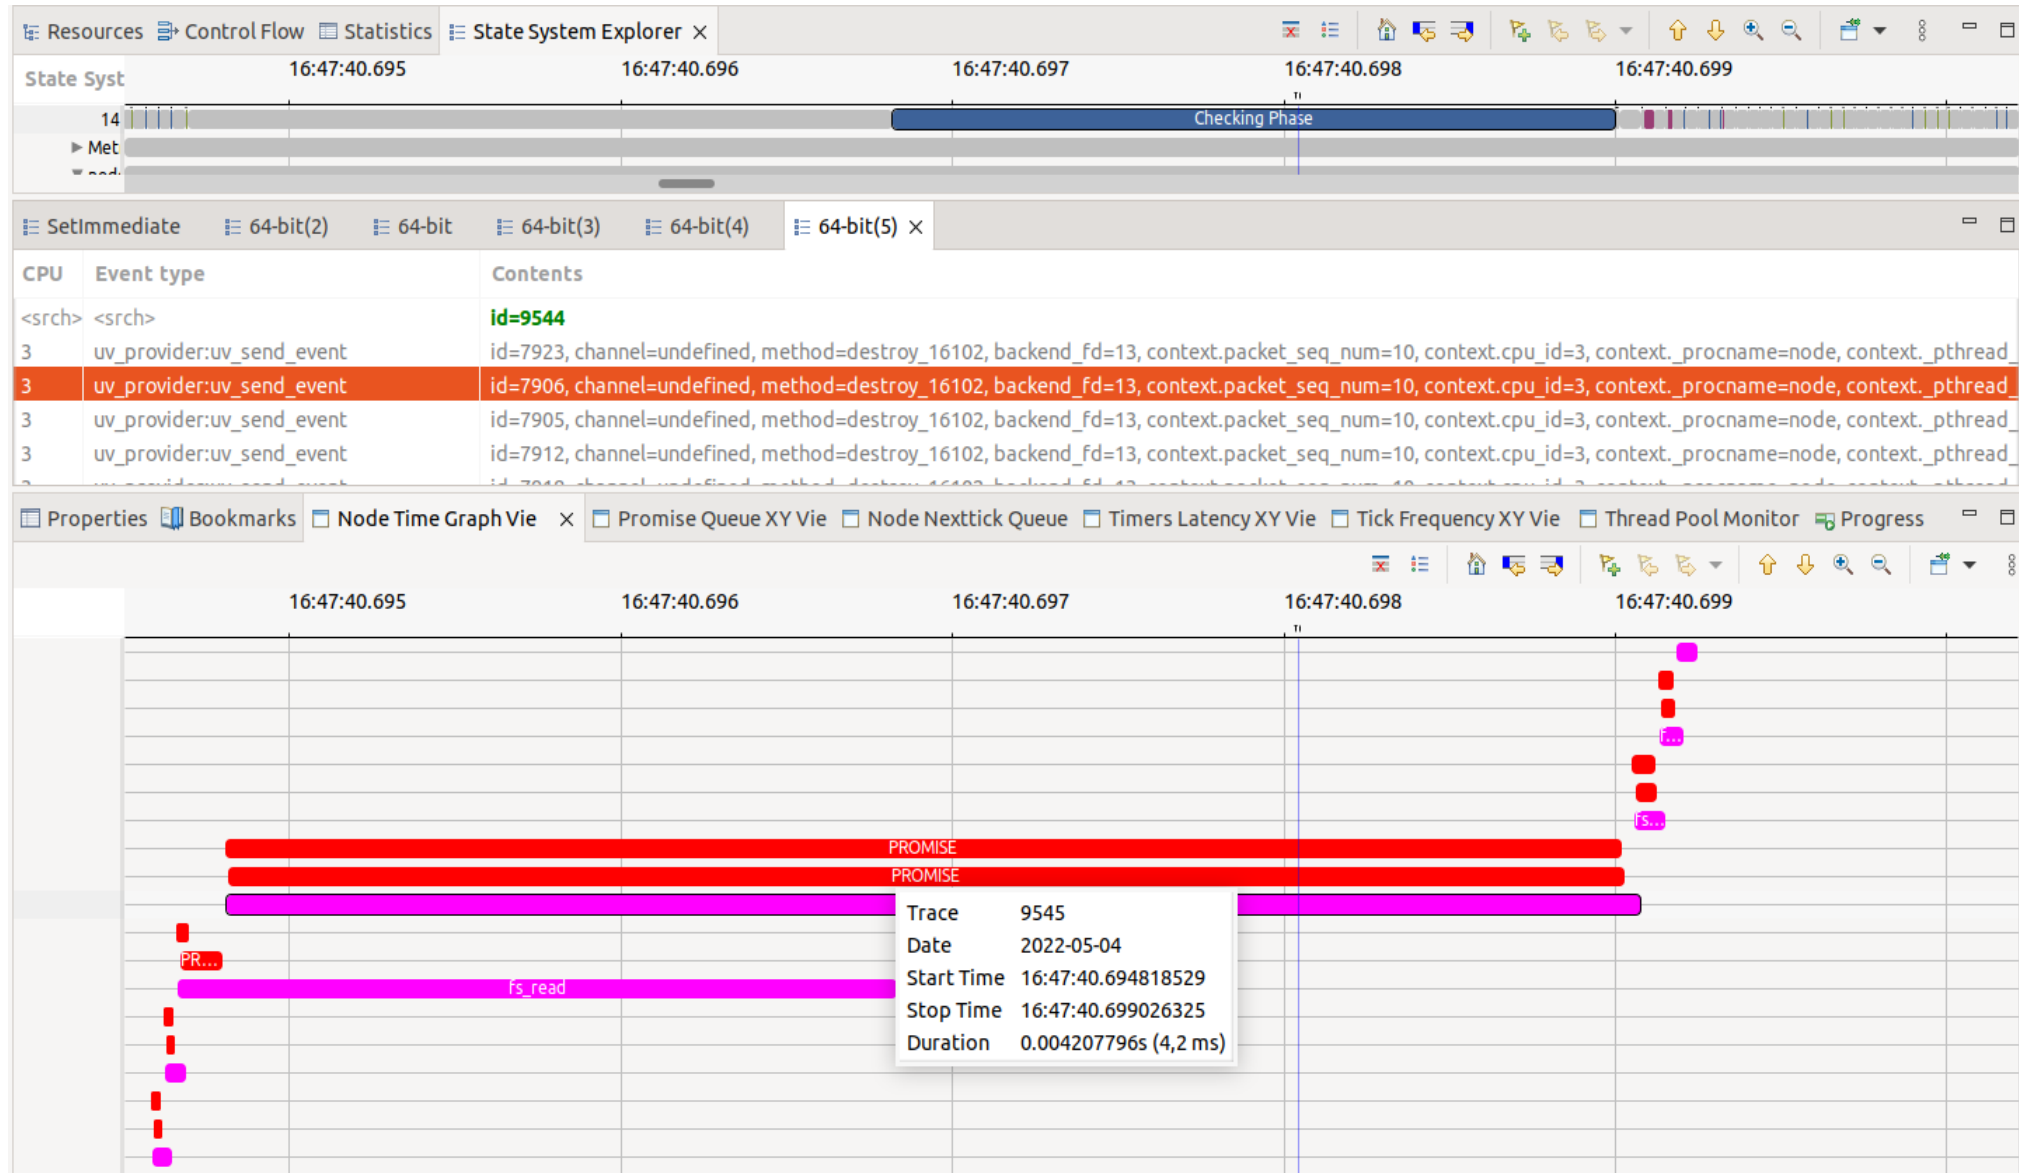
When contacted by A, the server calls the ***fs.promise.readFile()*** of the file system API

Use case 2

Resulting Flow

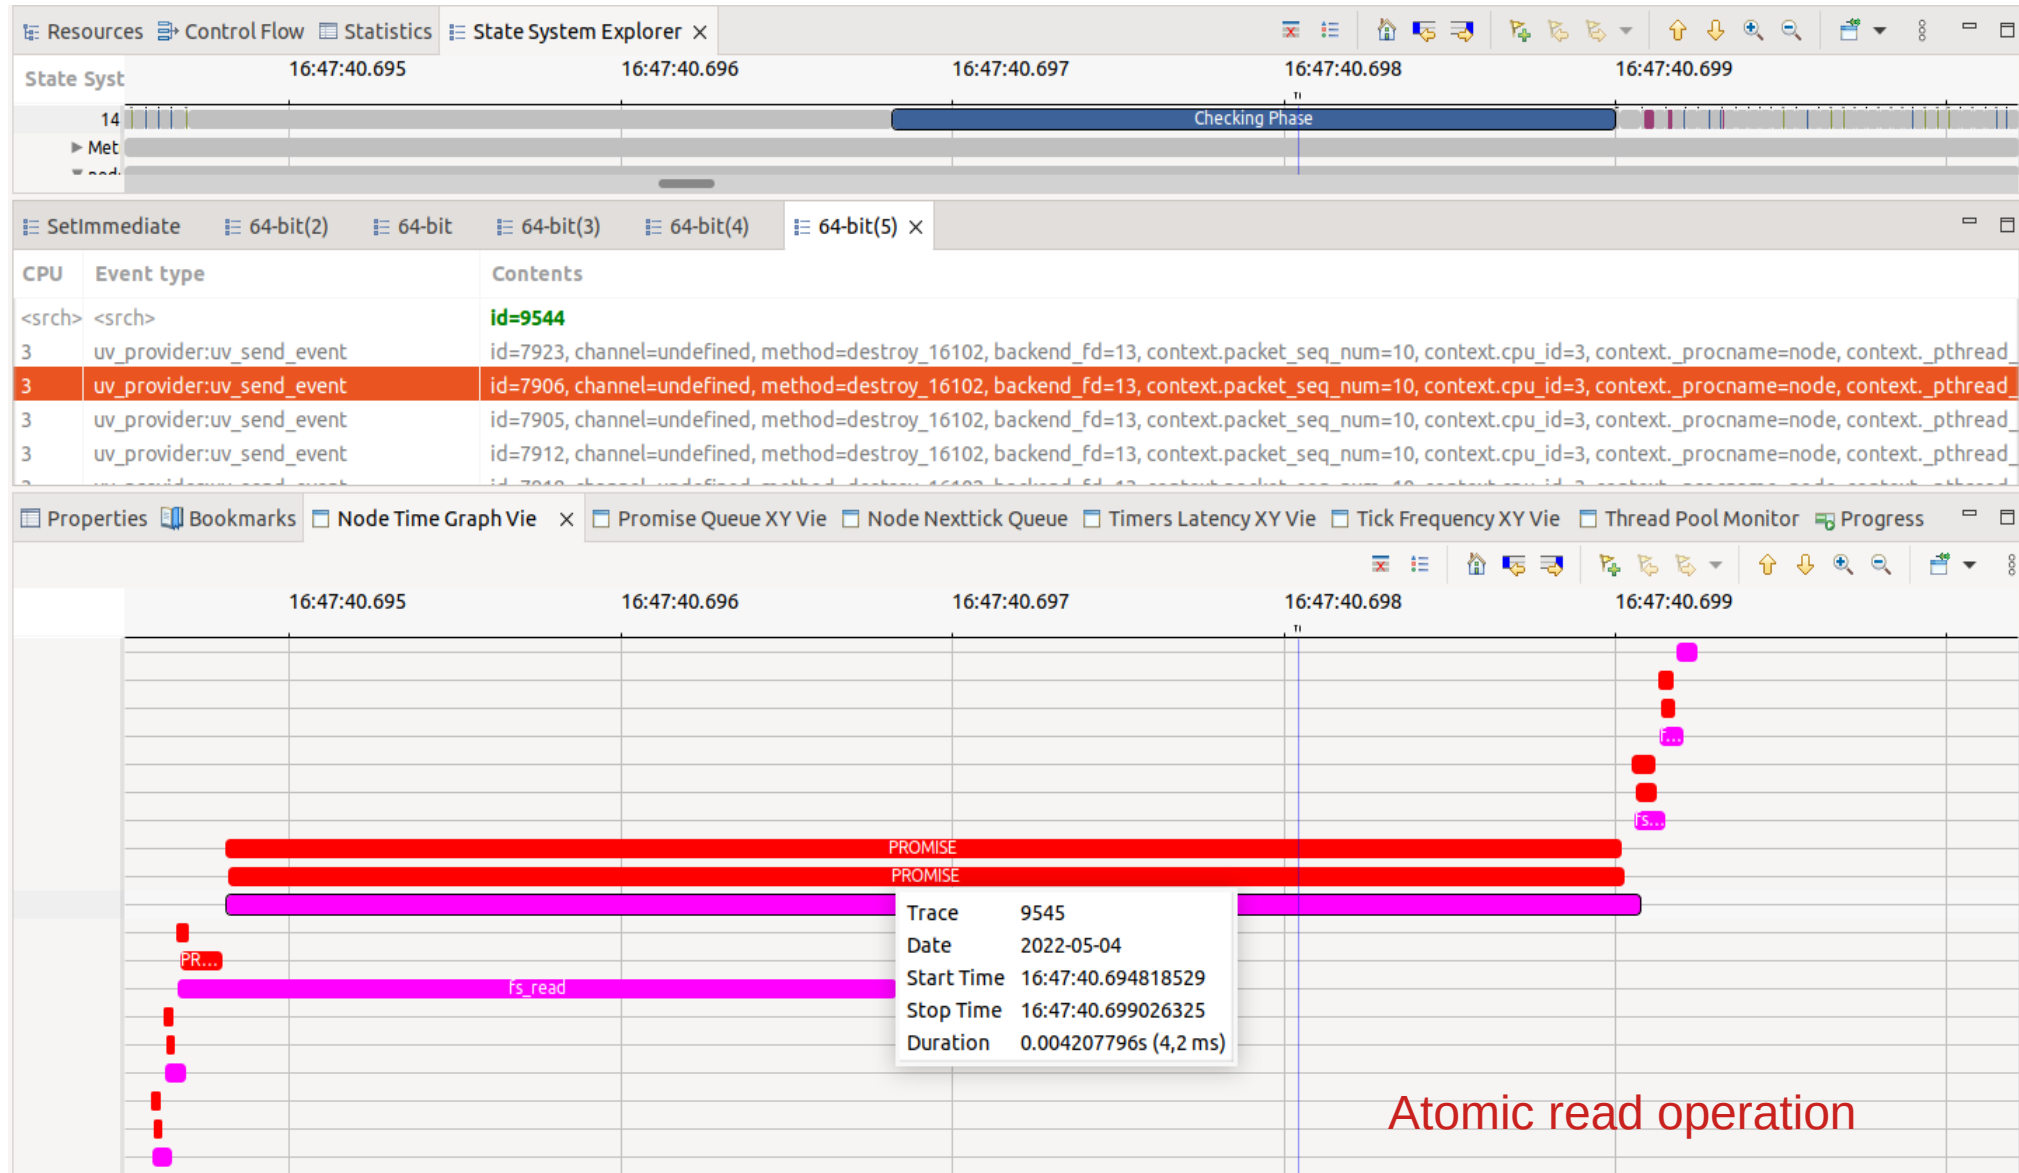
Use case 2

Resulting Flow



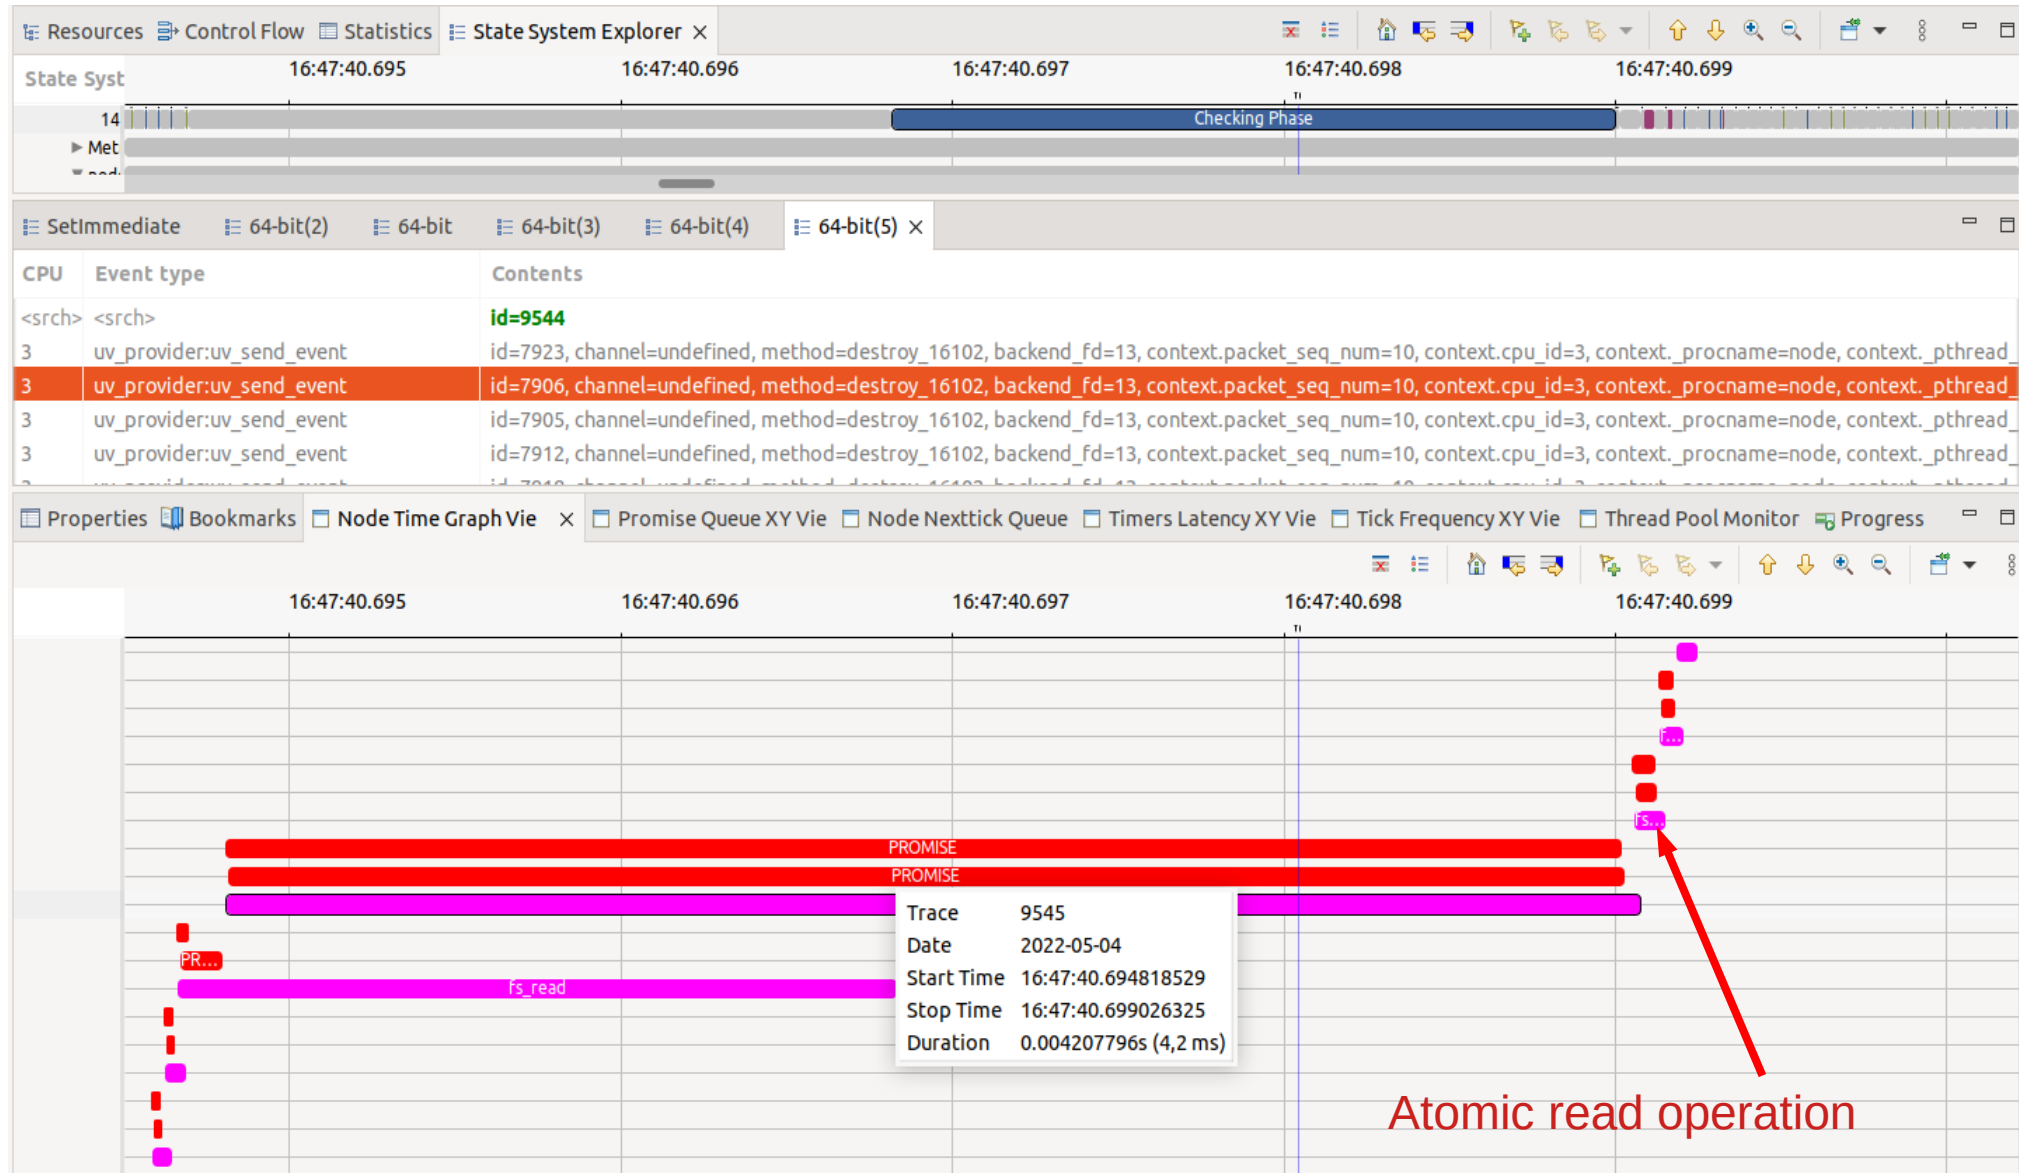
Use case 2

Resulting Flow



Use case 2

Resulting Flow

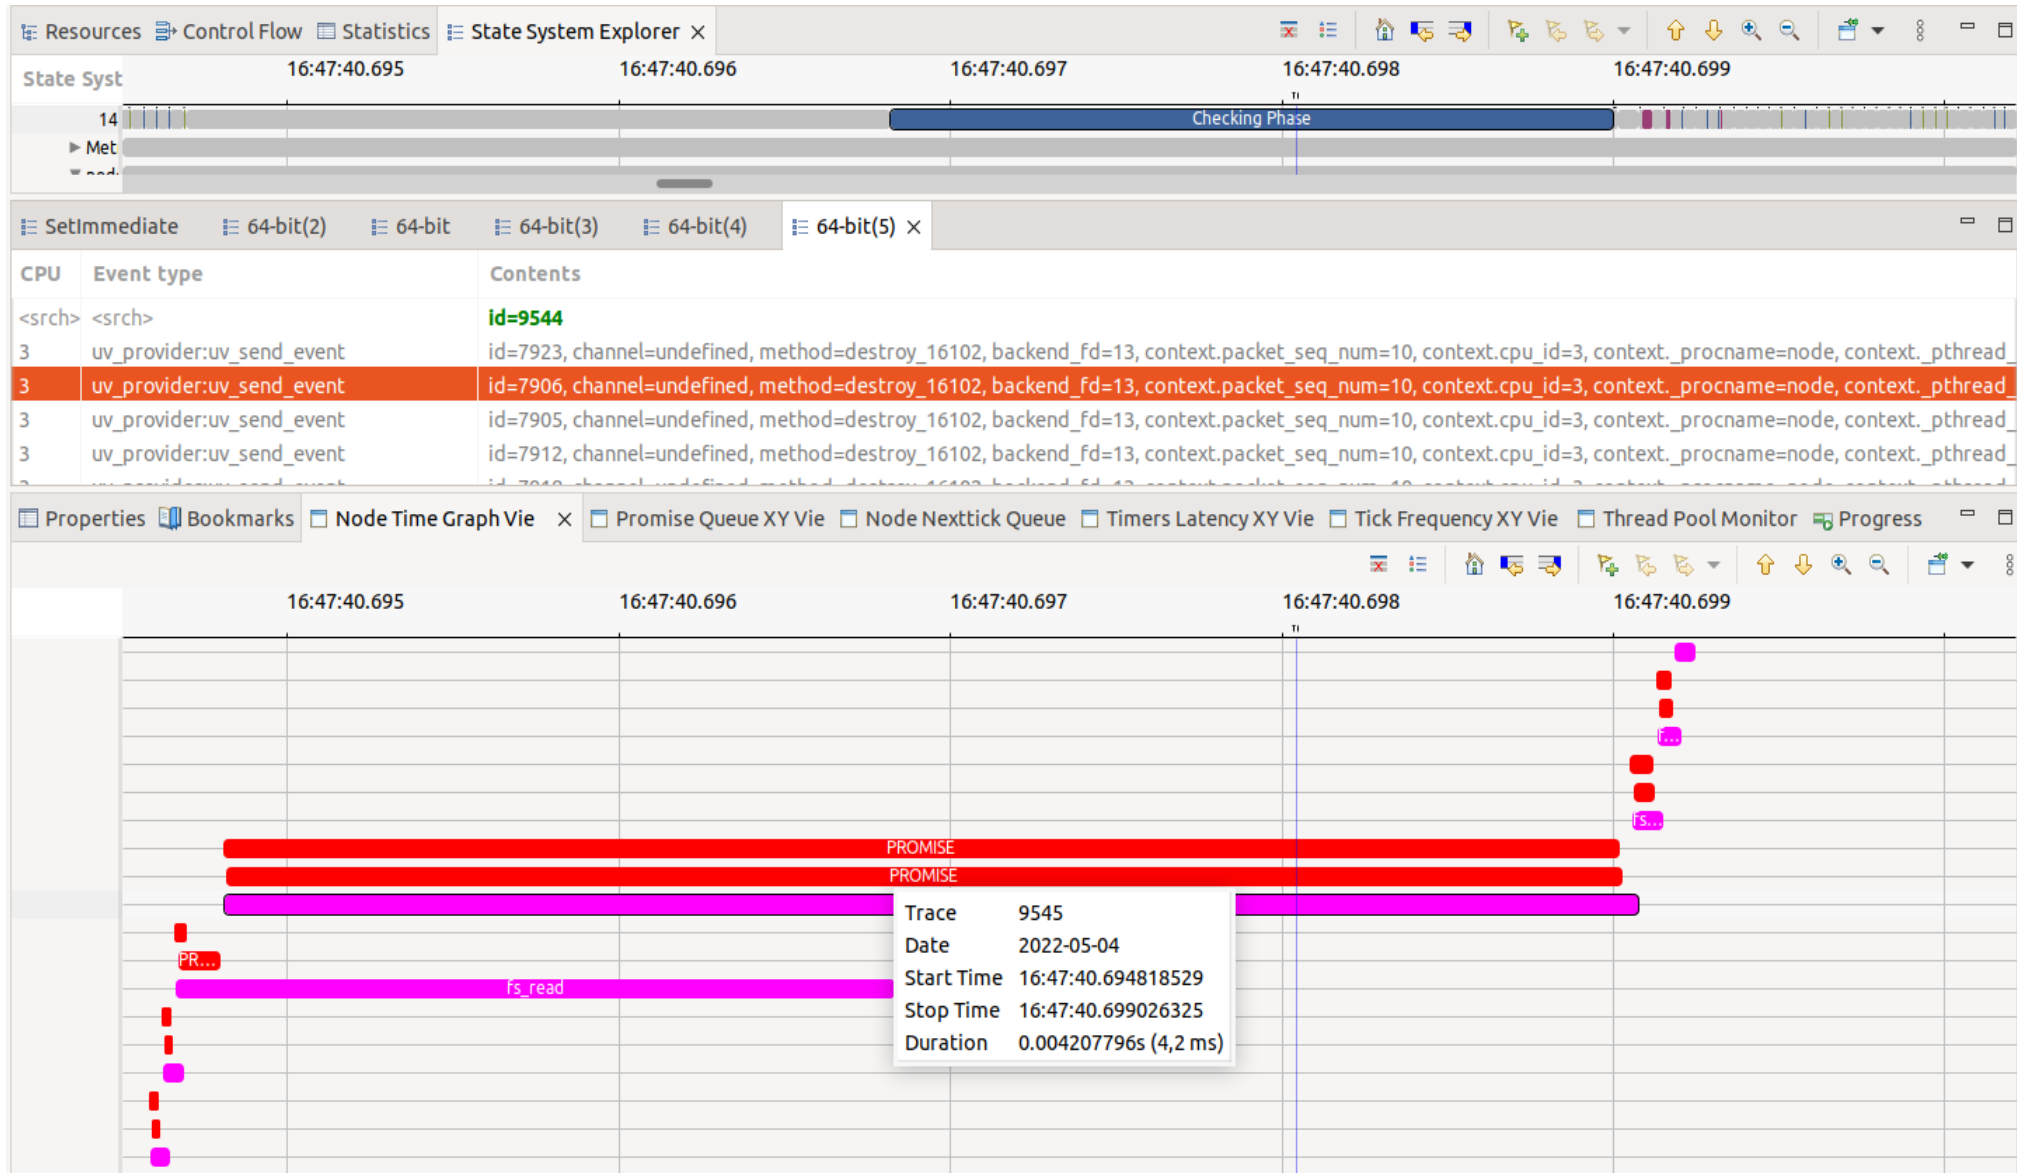


Use case 2

Results interpretation

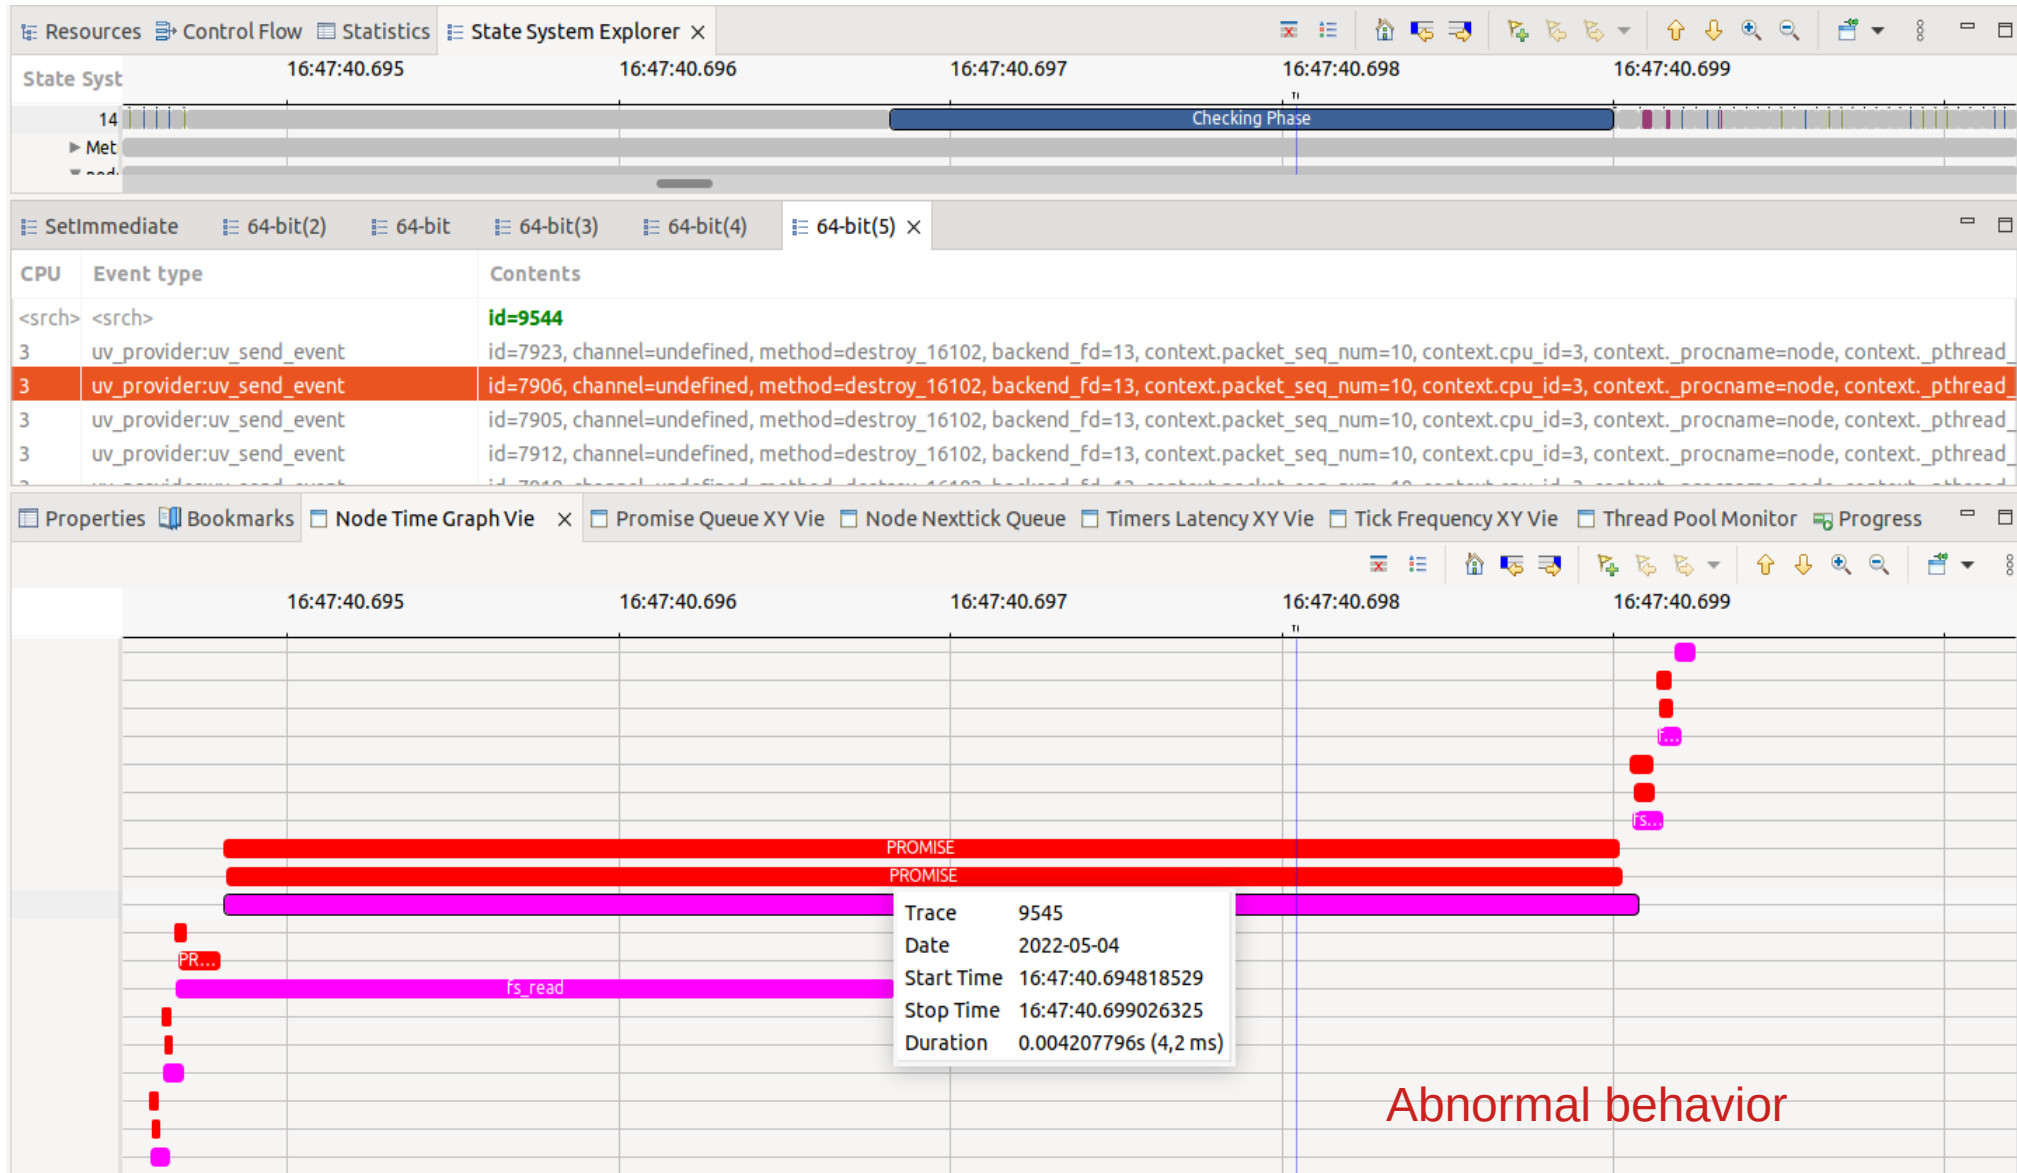
Use case 2

Results interpretation



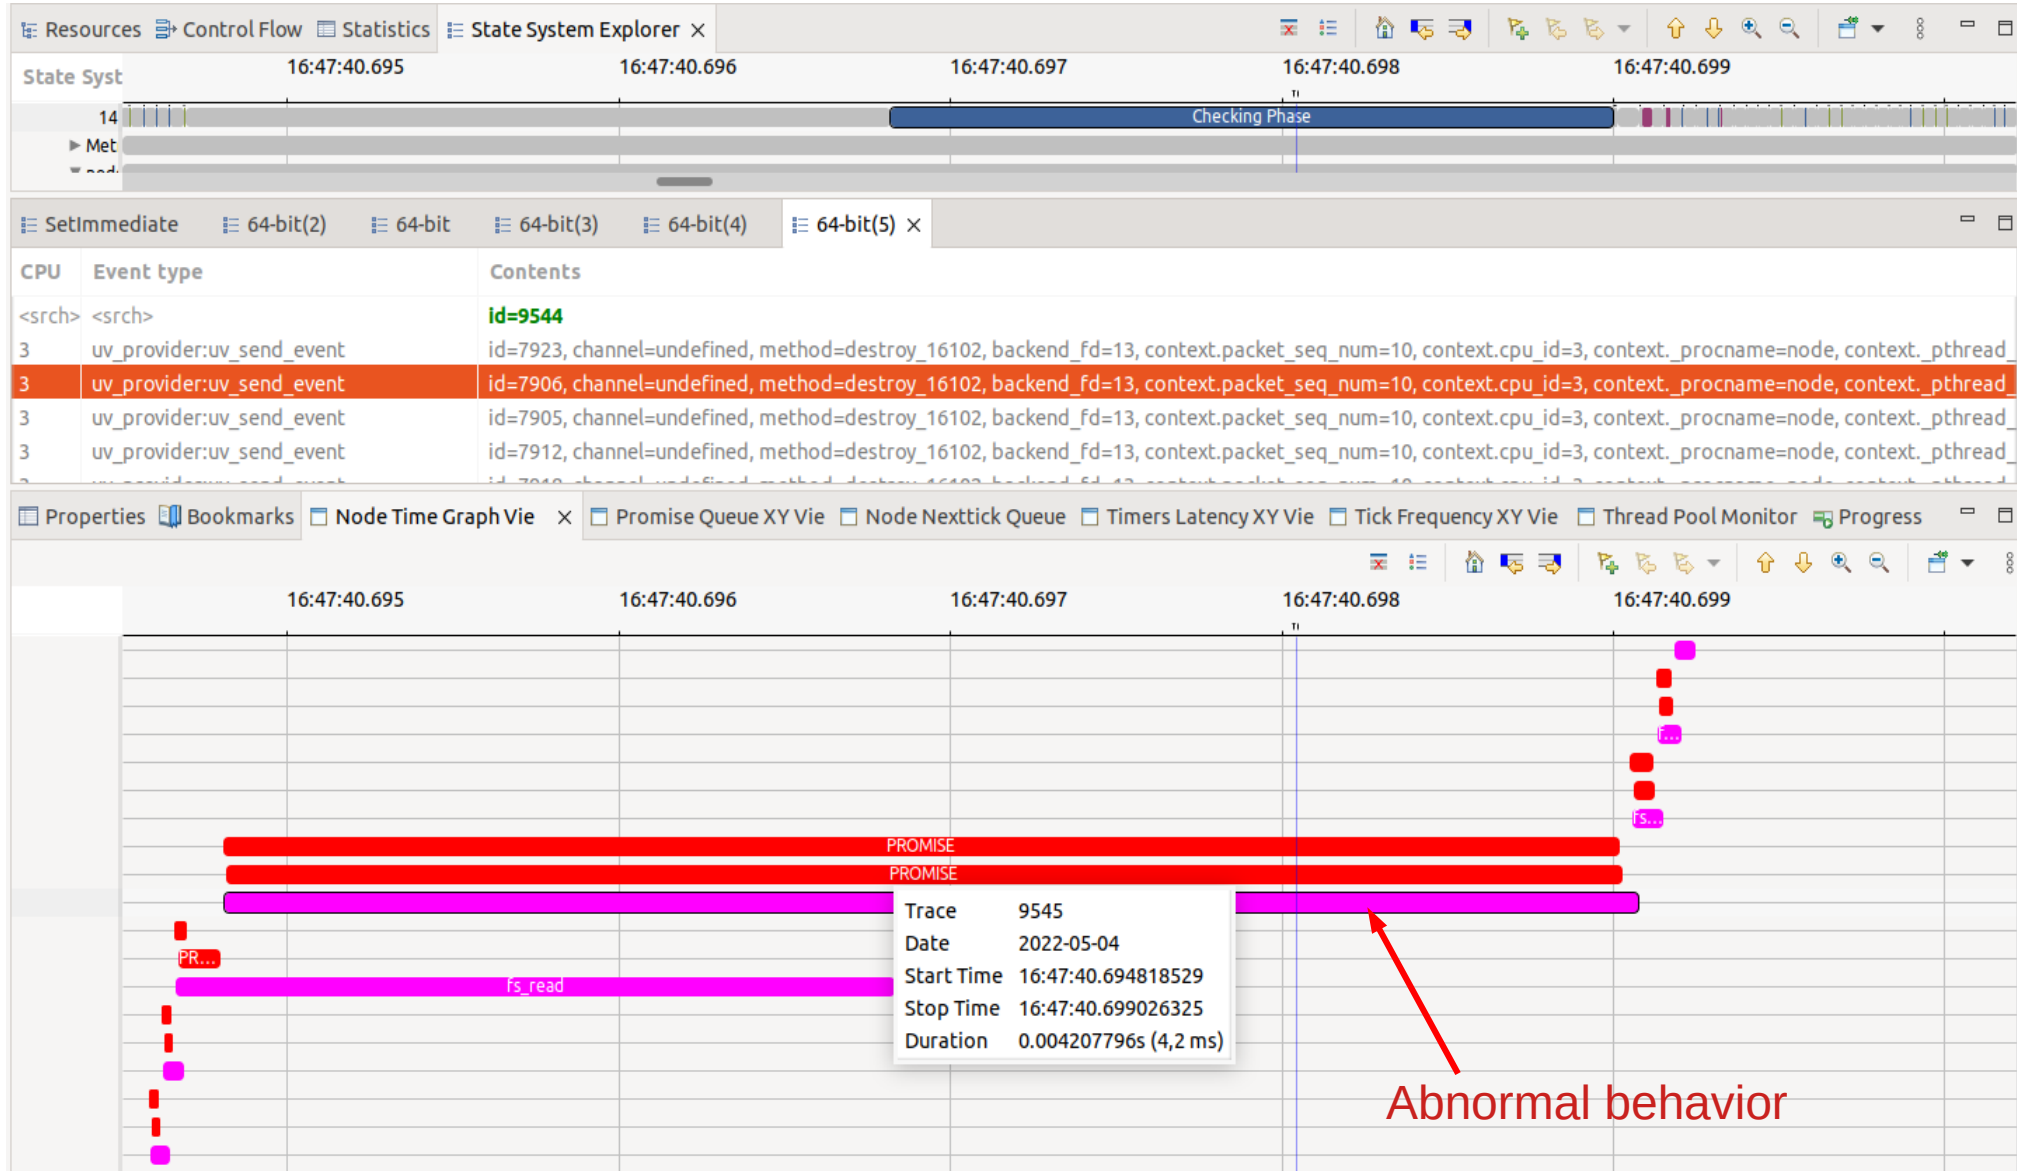
Use case 2

Results interpretation



Use case 2

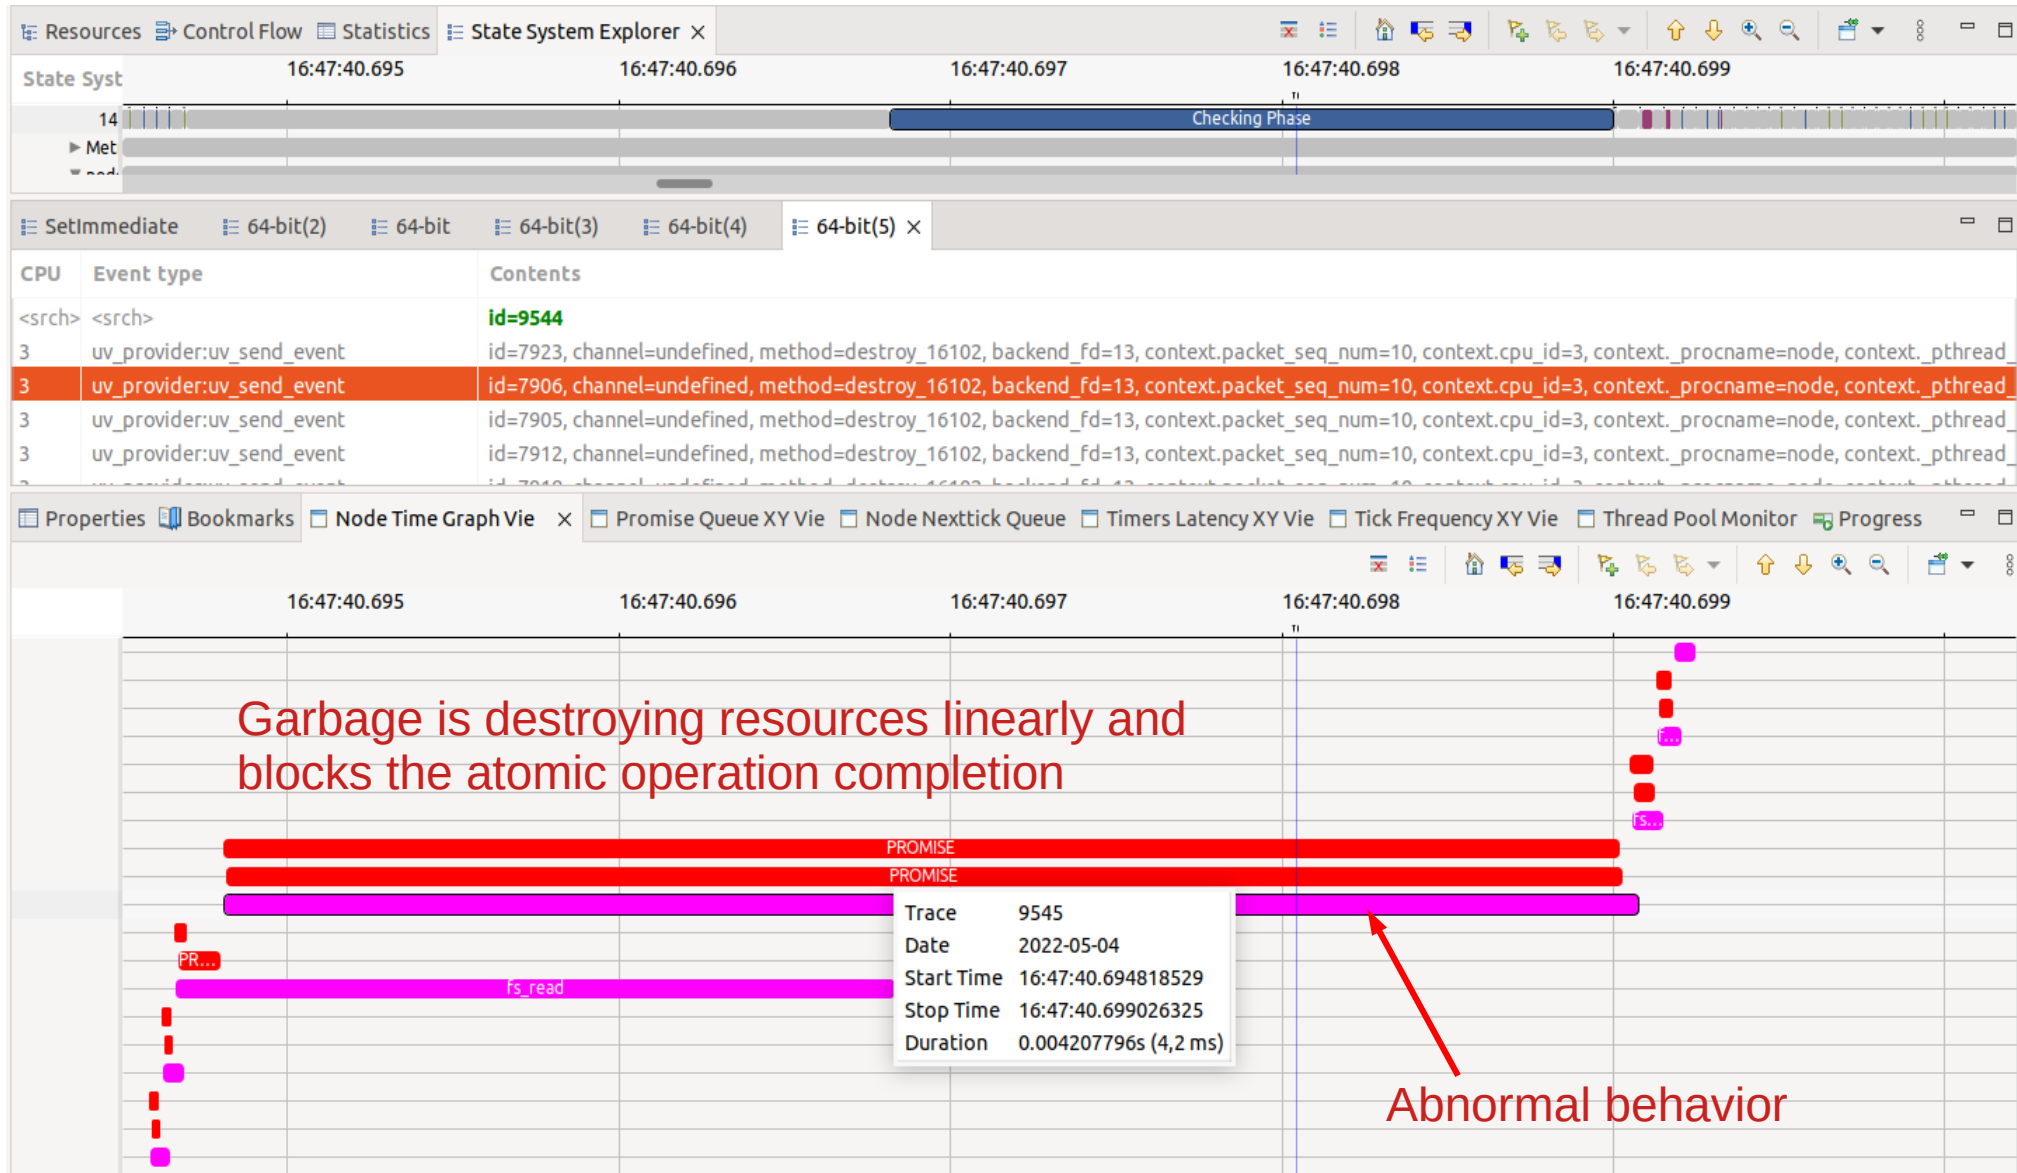
Results interpretation



Abnormal behavior

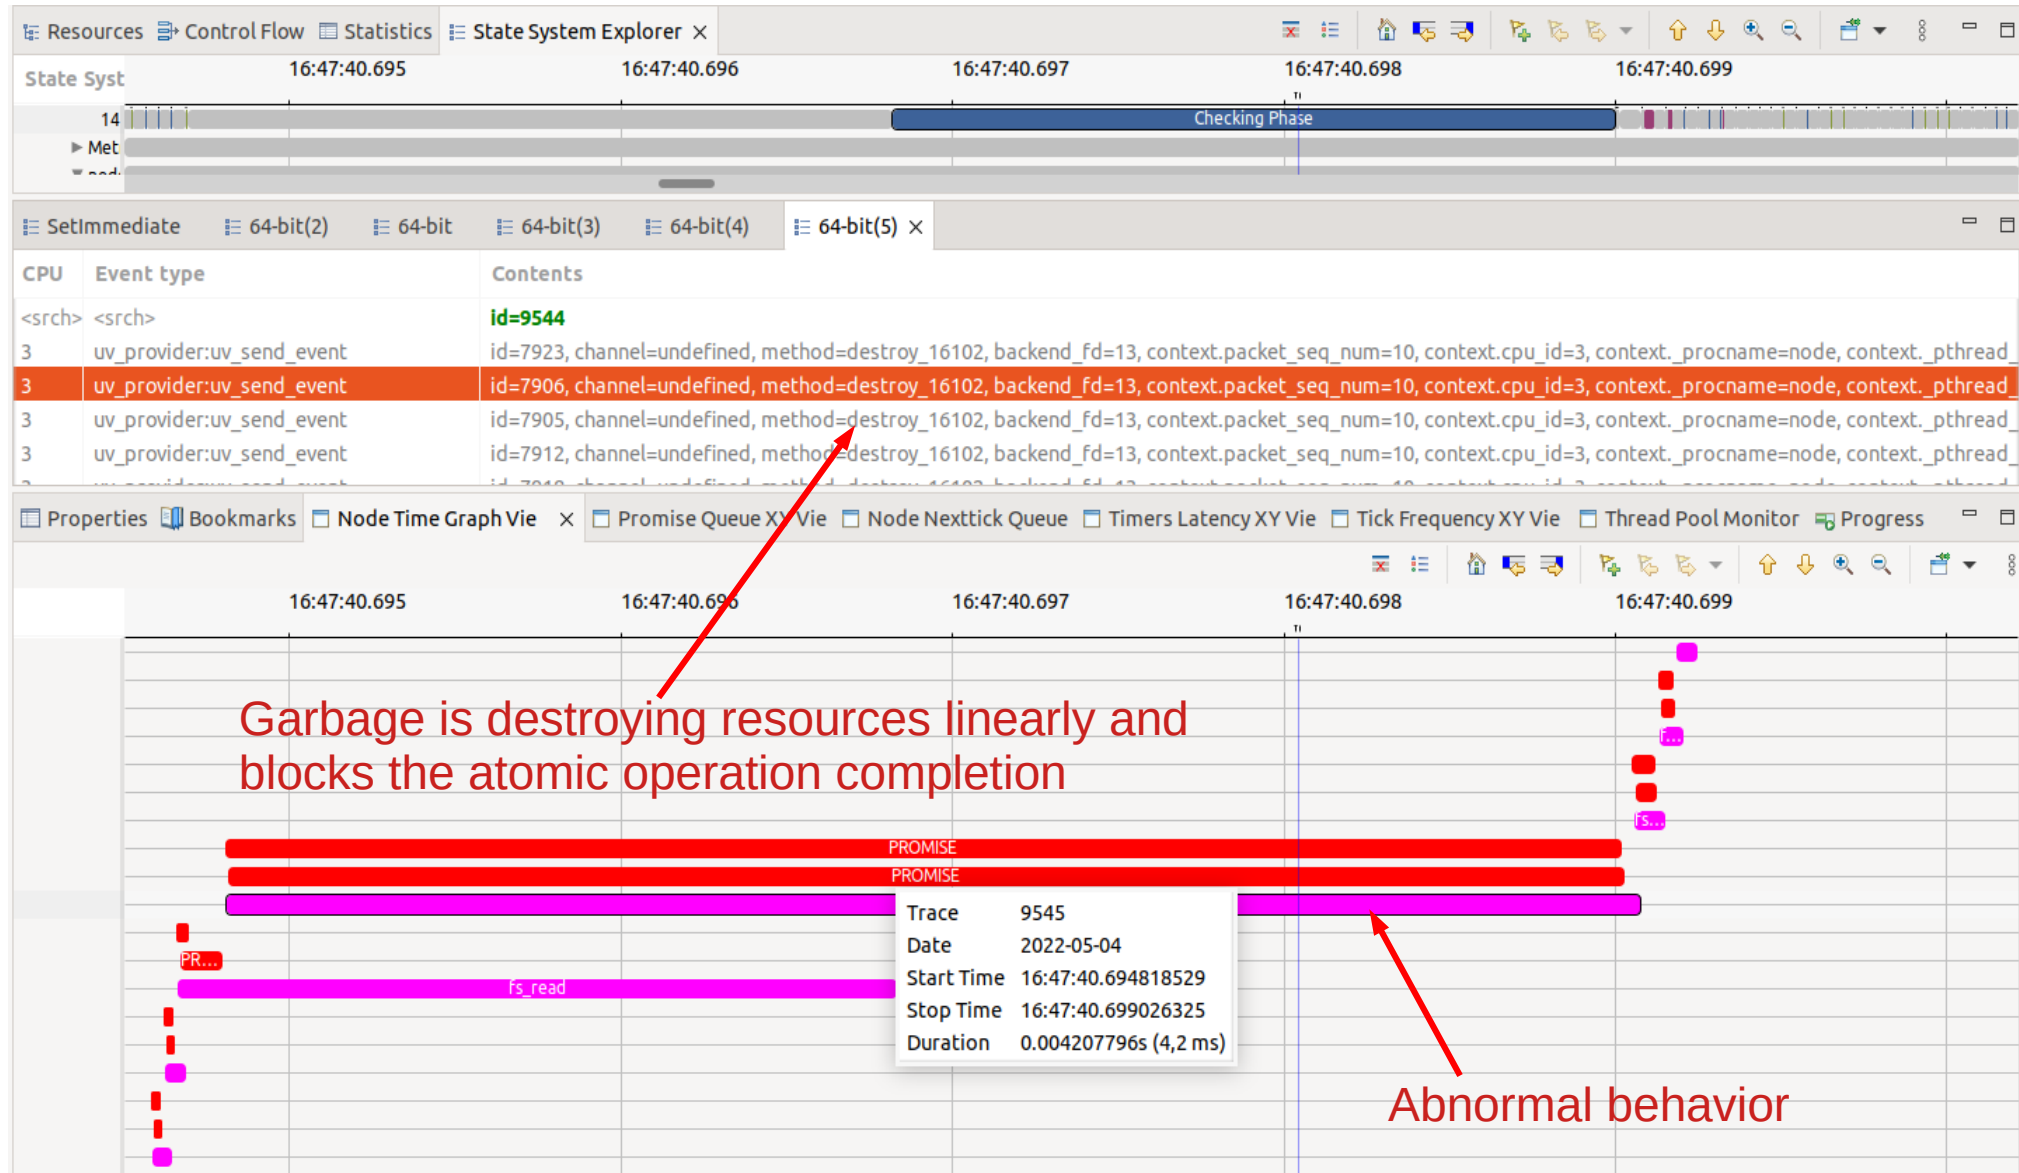
Use case 2

Results interpretation



Use case 2

Results interpretation



Use case 3

Race conditions detection

Use case 3

Race conditions detection

Atomic operation may access the same resource

Use case 3

Race conditions detection

Atomic operation may access the same resource

Two or more processes may access a shared resource

Use case 3

Race conditions detection

Atomic operation may access the same resource

Two or more processes may access a shared resource

The resource may be altered, locked, deleted or modified before the other accesses it

Test scenario

Use case 3

Race conditions detection

Atomic operation may access the same resource

Two or more processes may access a shared resource

The resource may be altered, locked, deleted or modified before the other accesses it

Test scenario

An express server contacted at the route /race

Use case 3

Race conditions detection

Atomic operation may access the same resource

Two or more processes may access a shared resource

The resource may be altered, locked, deleted or modified before the other accesses it

Test scenario

An express server contacted at the route /race

A client **A** contacting continuously the server on the route /race

Use case 3

Race conditions detection

Atomic operation may access the same resource

Two or more processes may access a shared resource

The resource may be altered, locked, deleted or modified before the other accesses it

Test scenario

An express server contacted at the route /race

A client **A** contacting continuously the server on the route /race

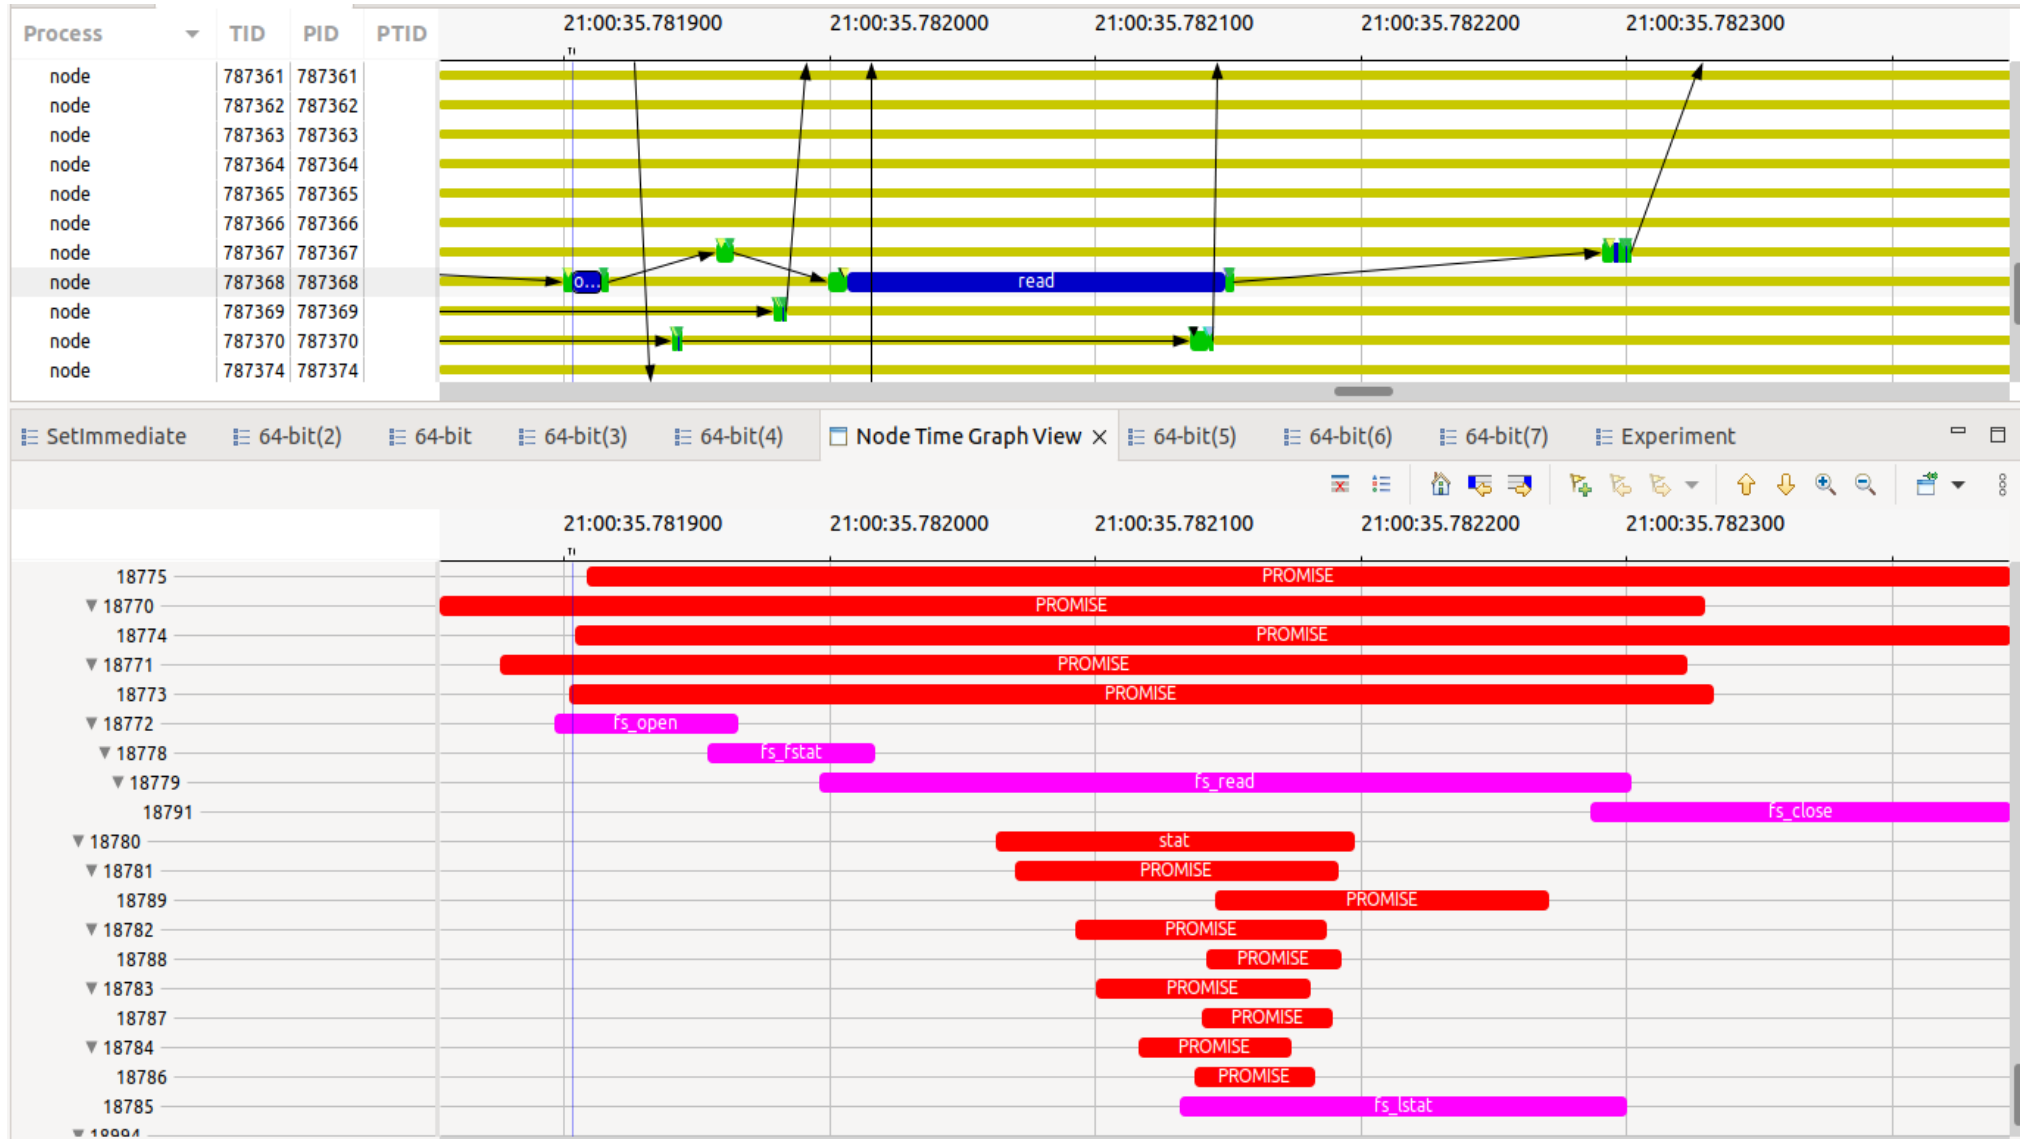
When contacted, the server calls a function that **creates** the file “test.txt”, then calls a function to **read** the file, and call another function to **delete** the file

Use case 3

Resulting flow

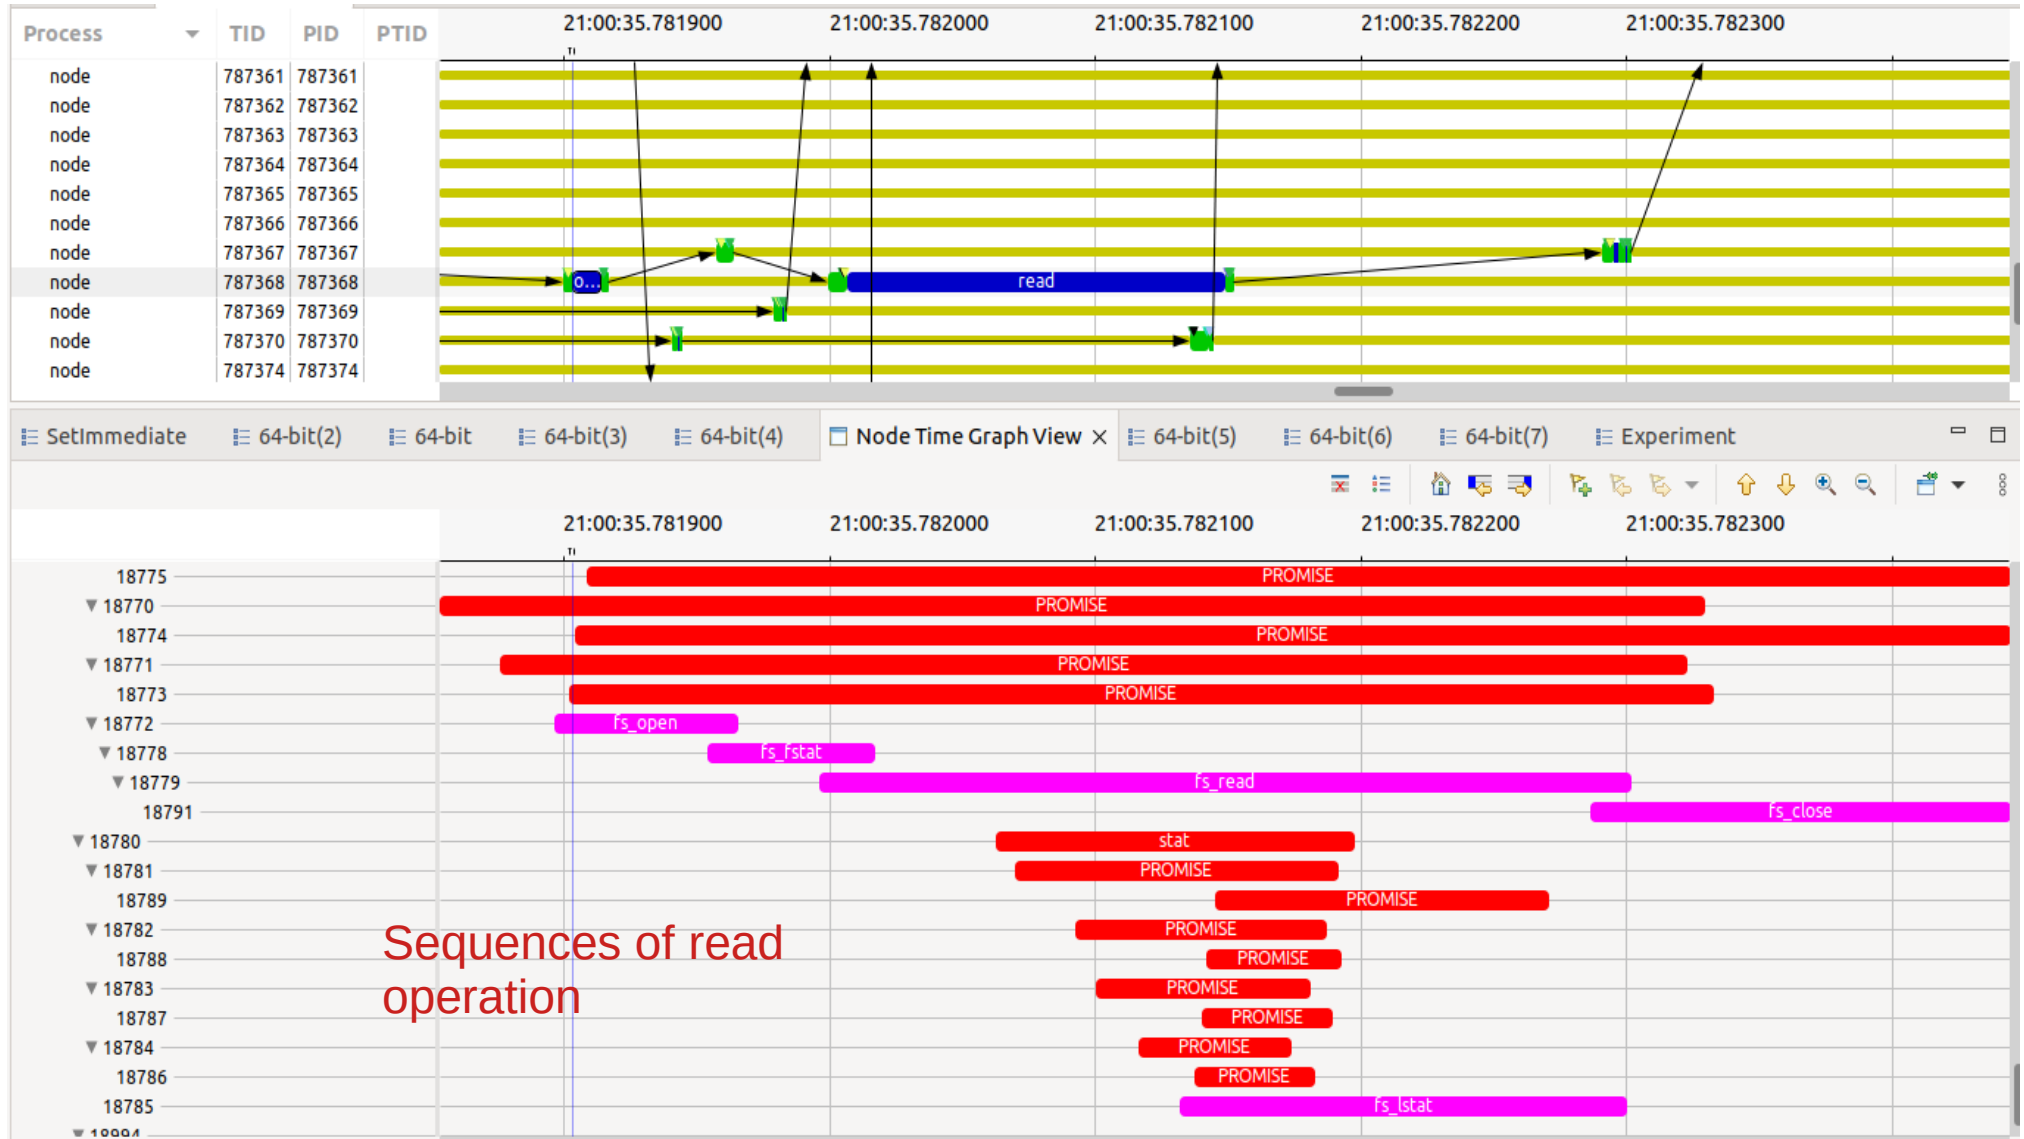
Use case 3

Resulting flow



Use case 3

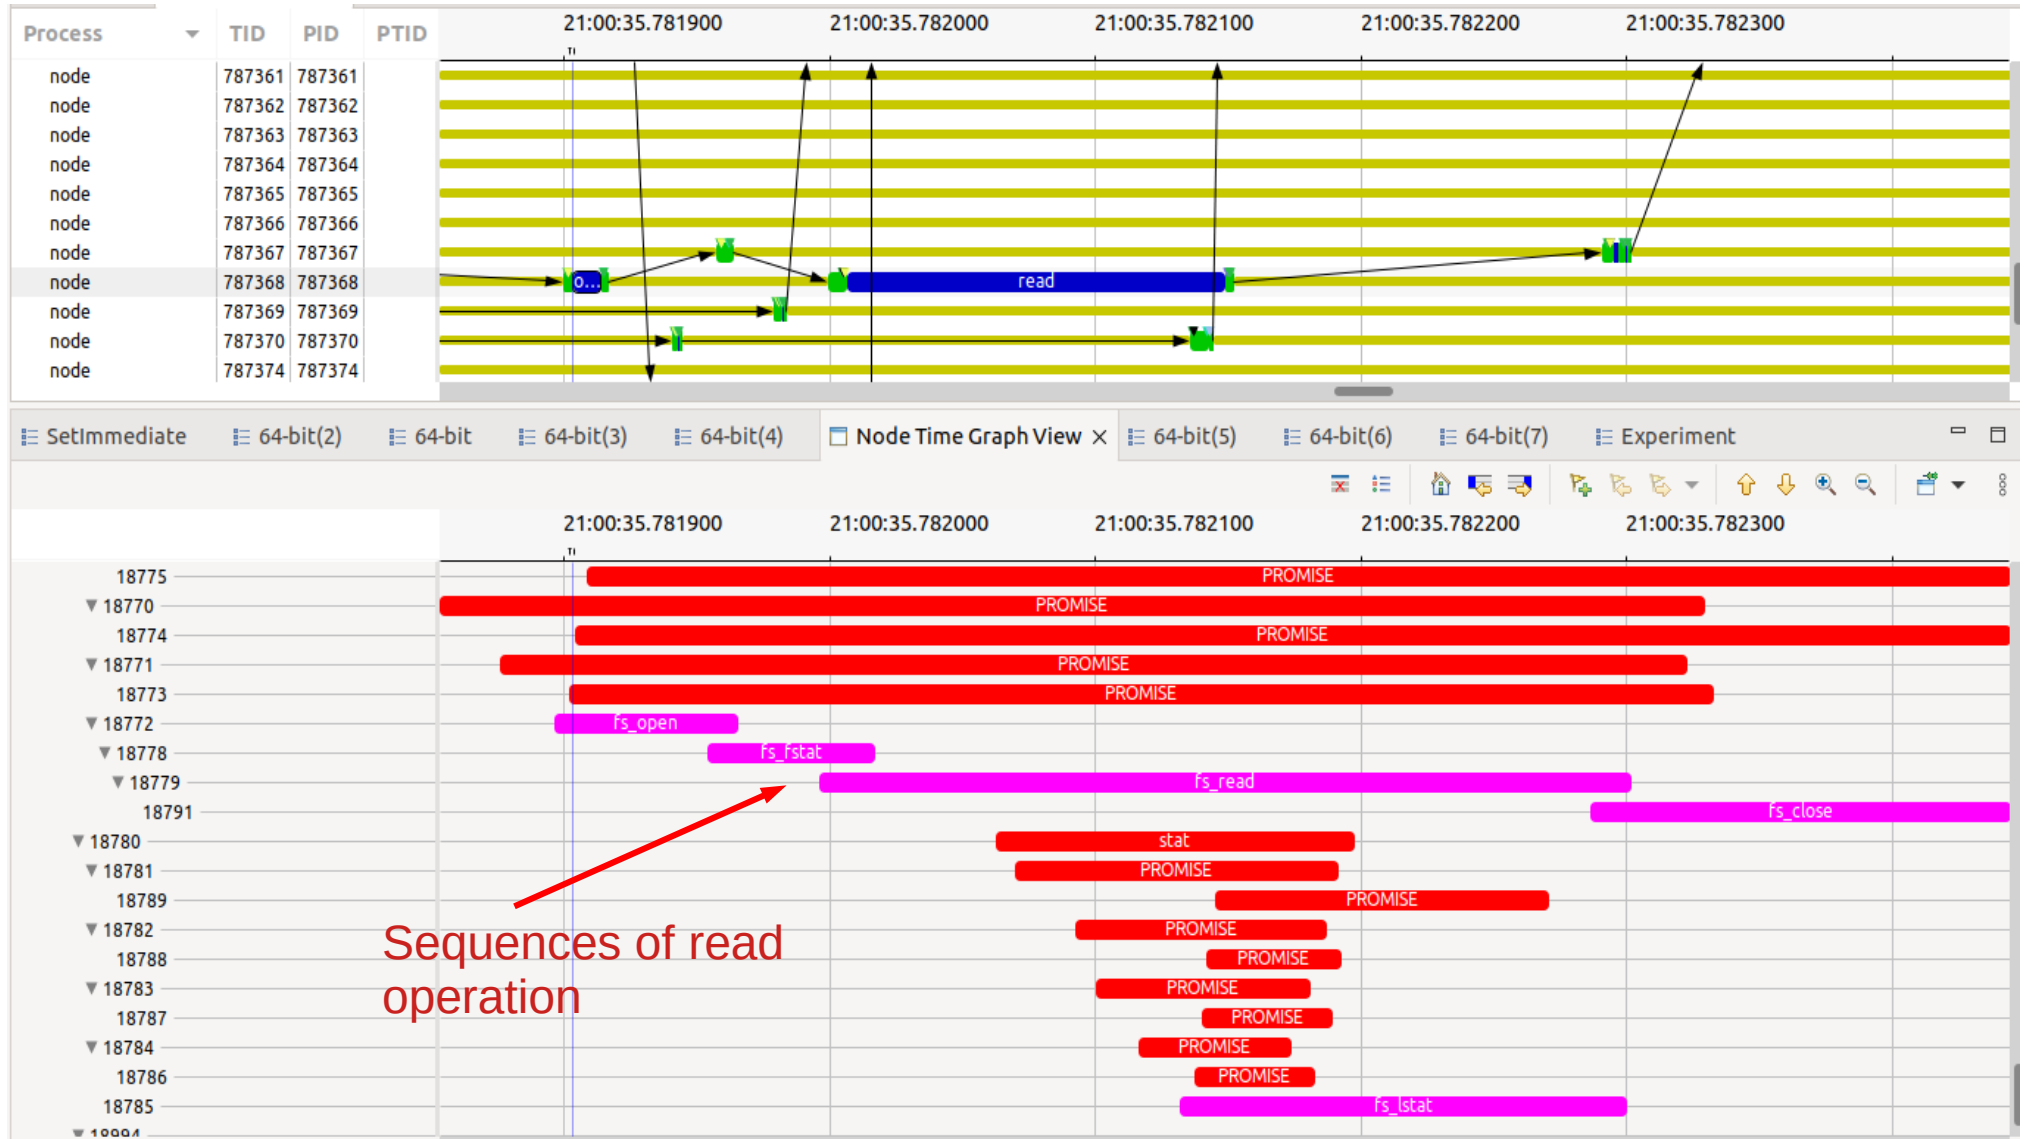
Resulting flow



Sequences of read operation

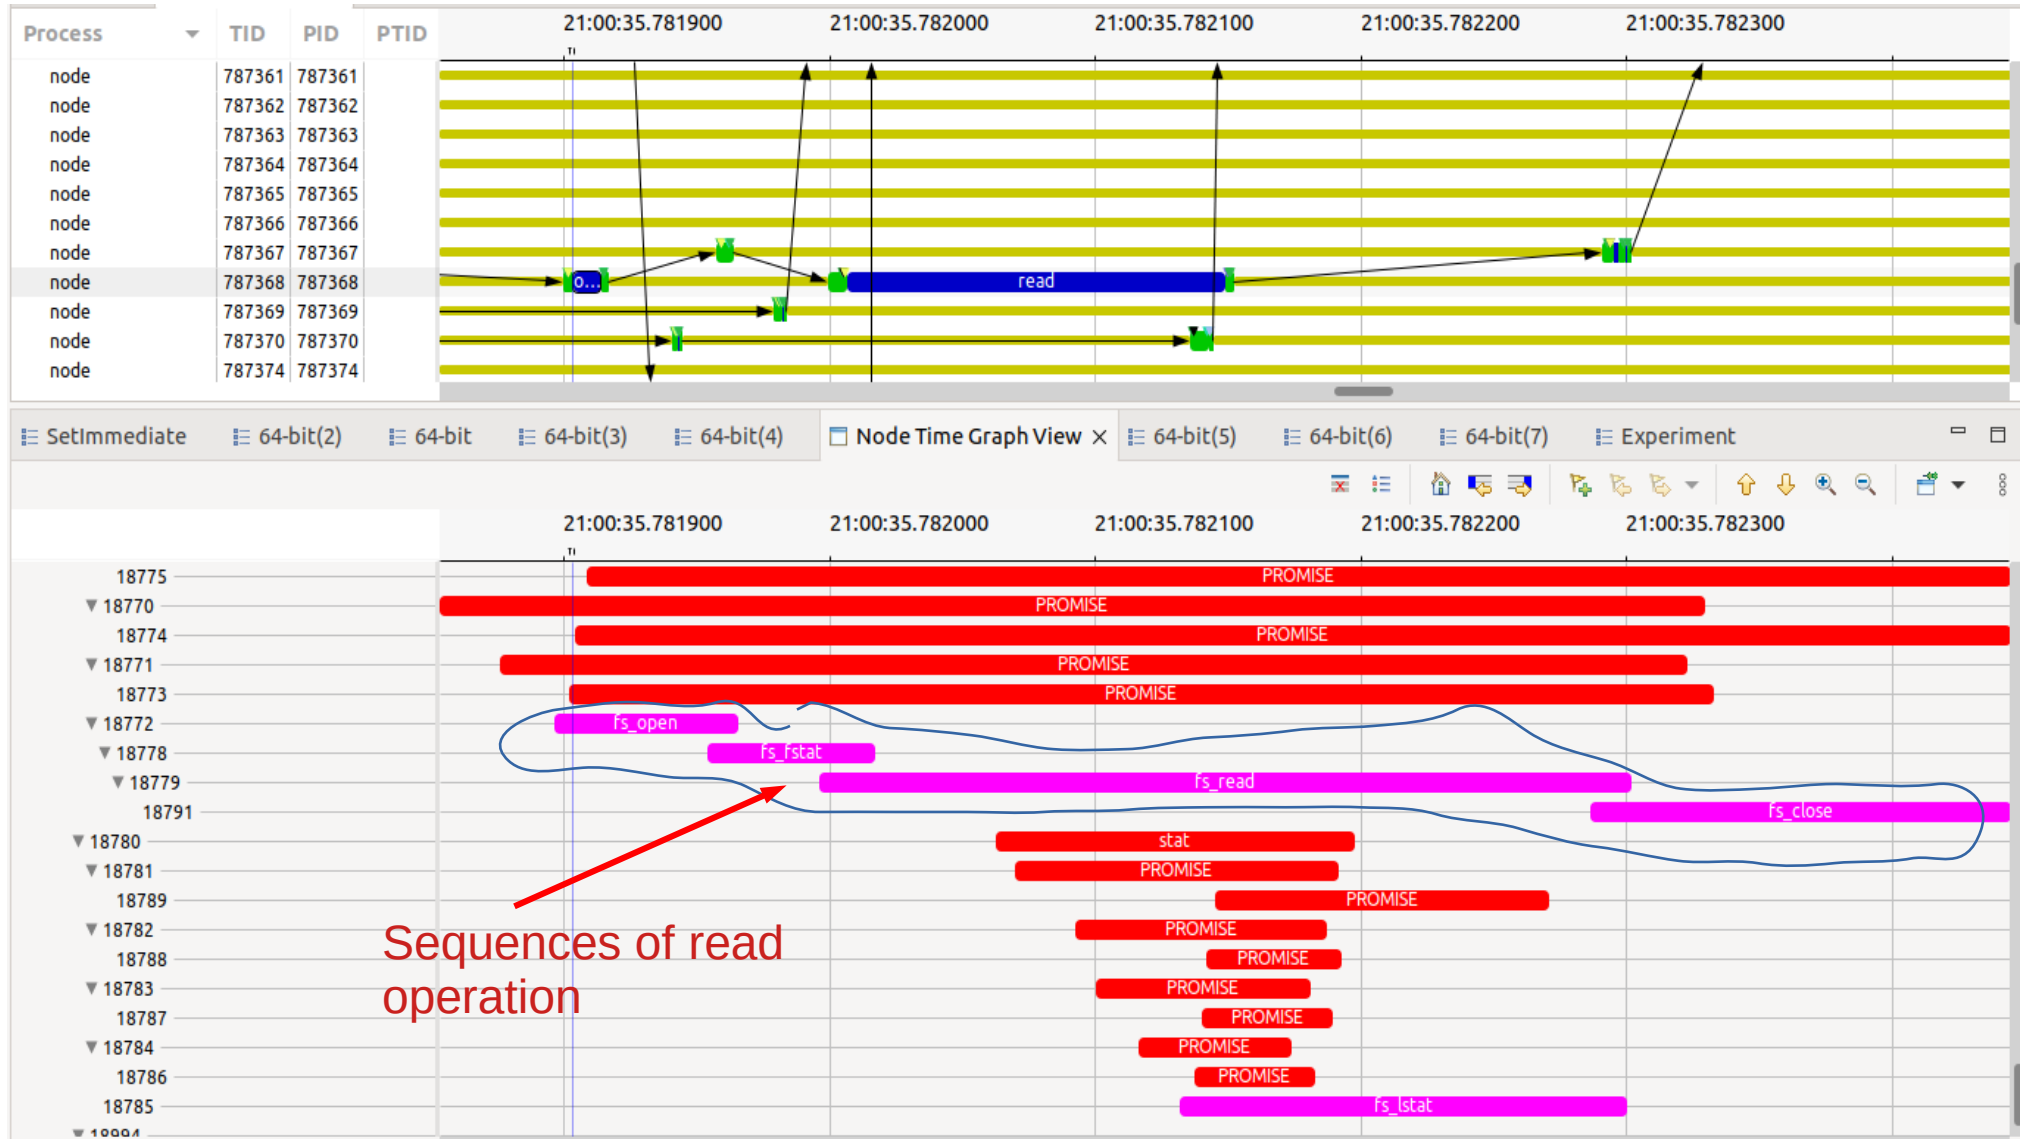
Use case 3

Resulting flow



Use case 3

Resulting flow

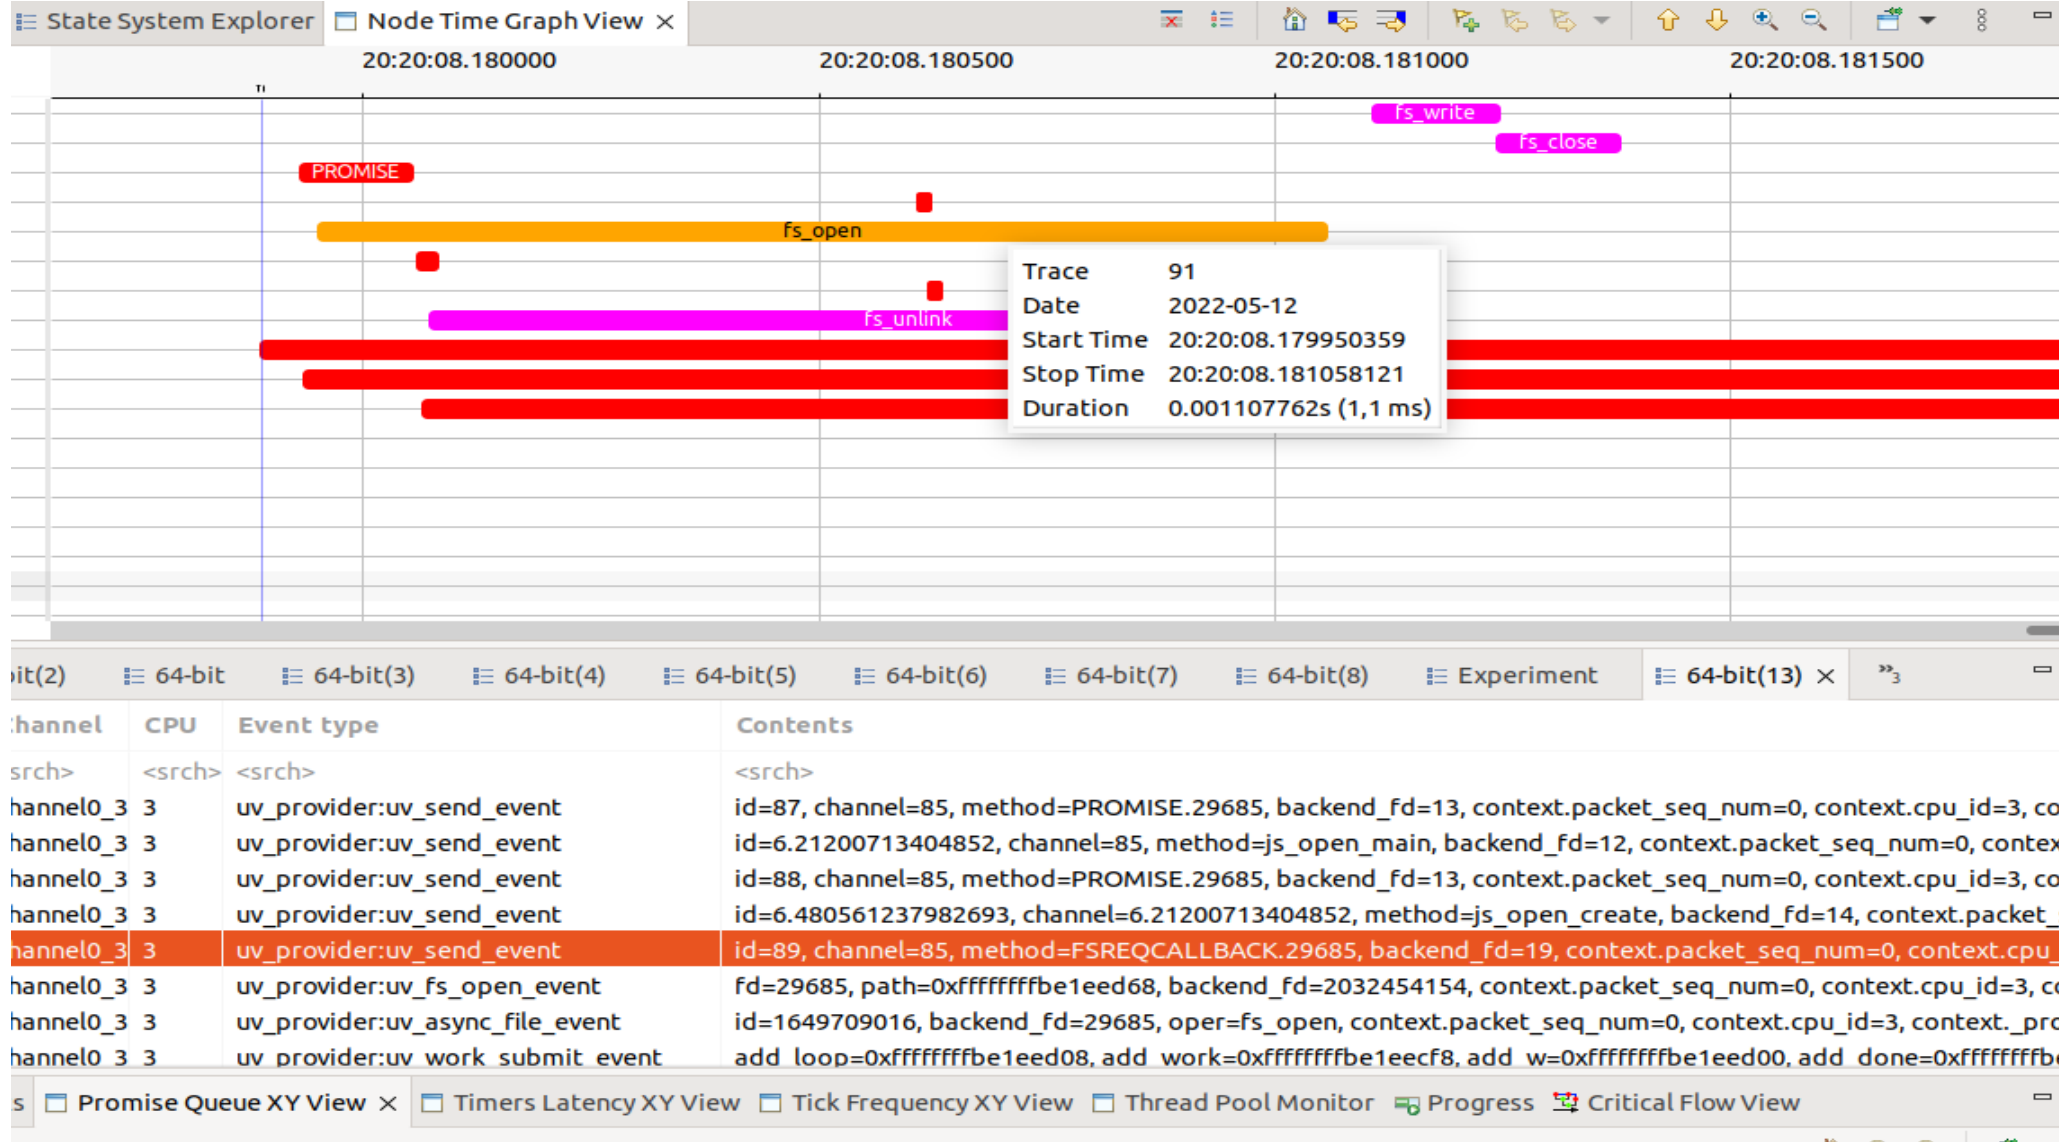


Use case 3

Results interpretation

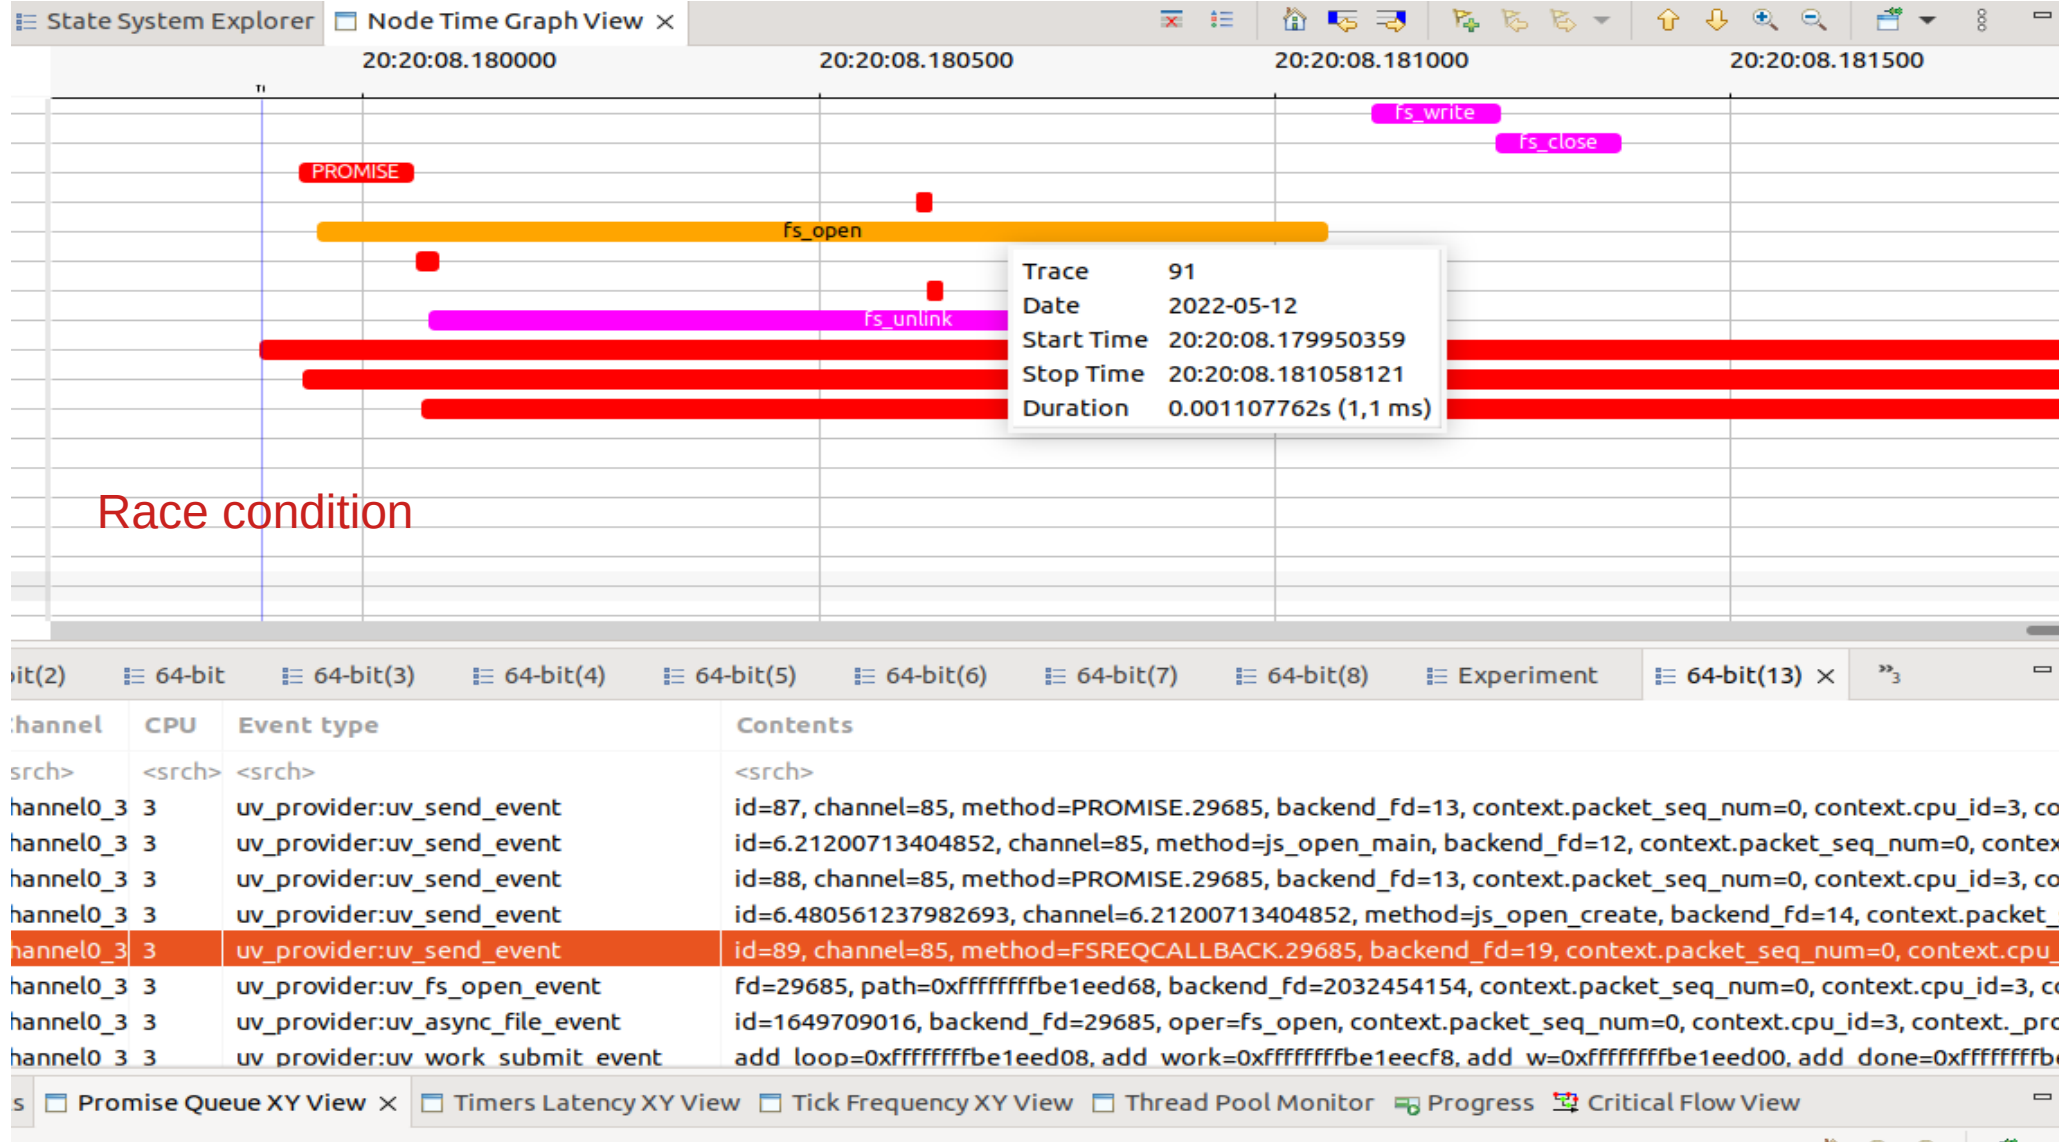
Use case 3

Results interpretation



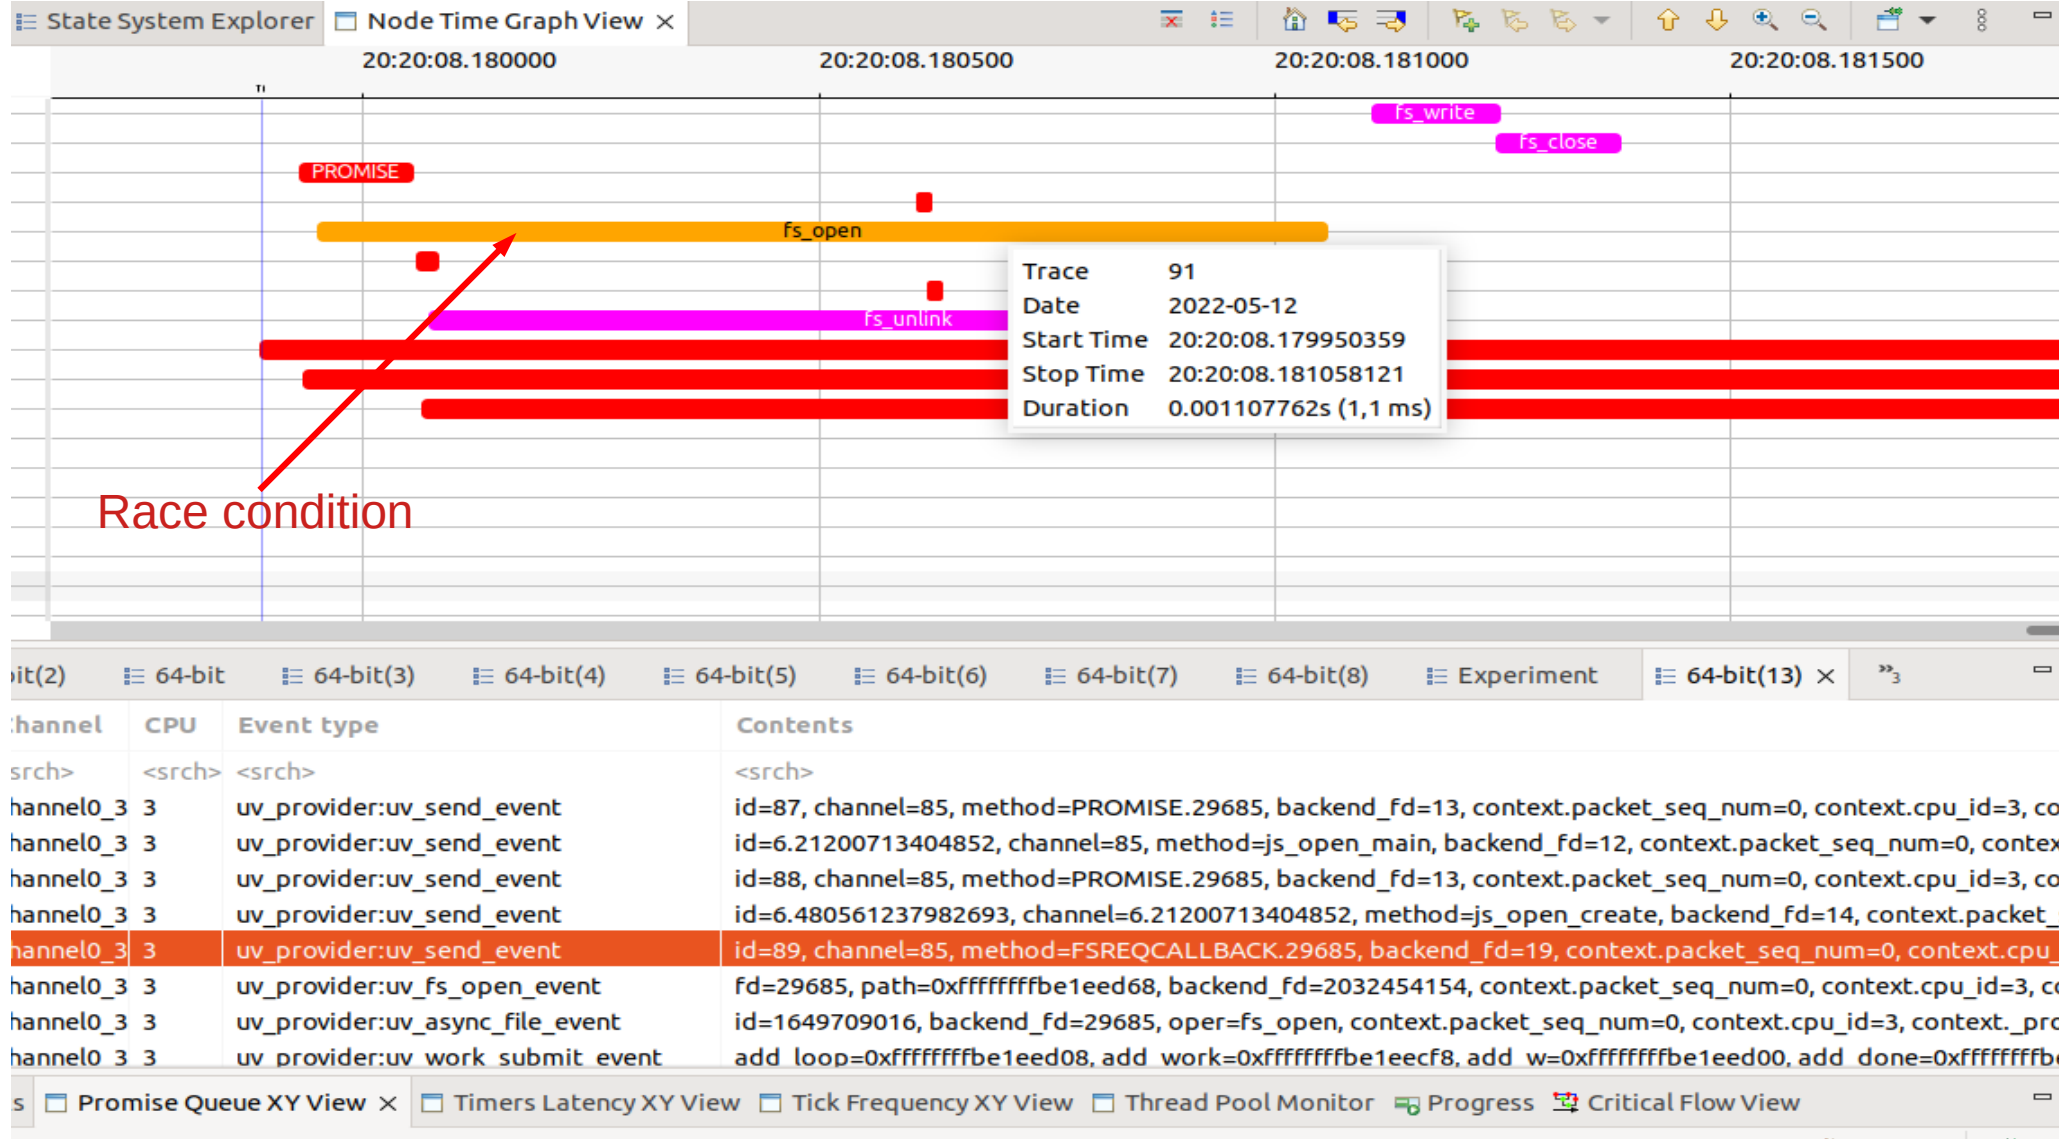
Use case 3

Results interpretation



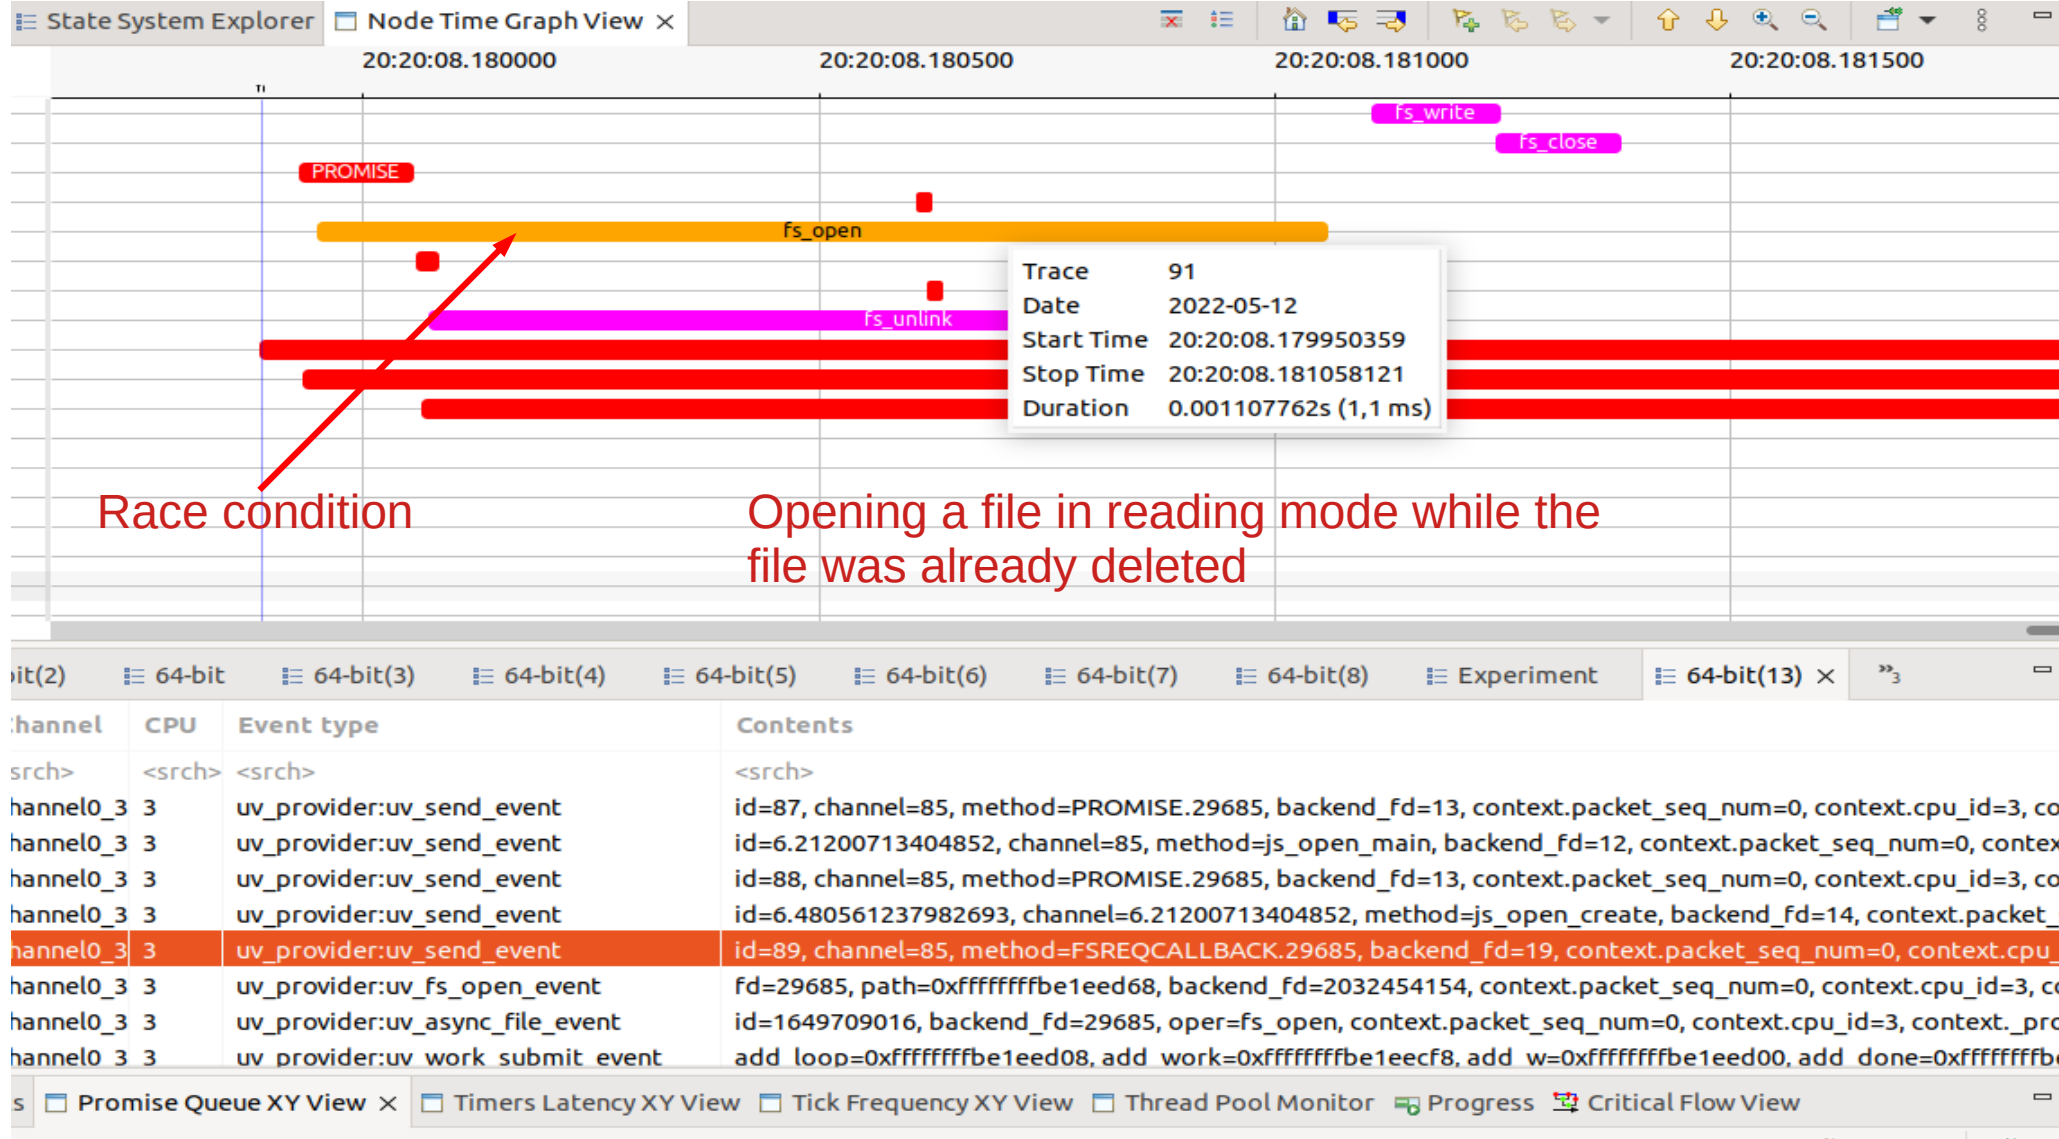
Use case 3

Results interpretation



Use case 3

Results interpretation



Conclusion

Conclusion

- NodeCompass a powerful tool for Node.js applications performance analysis

Conclusion

- NodeCompass a powerful tool for Node.js applications performance analysis
- Brings a high level of granularity in the Performance Analysis

Conclusion

- NodeCompass a powerful tool for Node.js applications performance analysis
- Brings a high level of granularity in the Performance Analysis
- Helps in understanding the application flow

Conclusion

- NodeCompass a powerful tool for Node.js applications performance analysis
- Brings a high level of granularity in the Performance Analysis
- Helps in understanding the application flow
- Allows pinpointing bottlenecks, bugs, errors, race conditions

Conclusion

- NodeCompass a powerful tool for Node.js applications performance analysis
- Brings a high level of granularity in the Performance Analysis
- Helps in understanding the application flow
- Allows pinpointing bottlenecks, bugs, errors, race conditions
- Can be used to instrument functions in the way distributed tracers work

Bibliography

- [1] I. Beschastnikh, P. Wang, Y. Brun, M. D. Ernst, Debugging distributed systems, ACM-Queue (2015).
- [2] J. Hoglund, An analysis of a distributed tracing systems effect on performance. jaeger and opentracing api, UMEA University (2020).
- [3] S. Tilkov, S. Vinoski, Node.js: Using javascript to build high-performance network programs, IEEE INTERNET COMPUTING (2010).
- [4] Cloud desktop ide platform.
URL <https://kubernetes.io/fr/>
- [5] Visual studio code.
URL <https://code.visualstudio.com/>
- [6] Y. Geng, S. Liu, Z. Yin, A. Naik, B. Prabhakar, M. Rosenblum, A. Vahdat, Exploiting a natural network effect for scalable, fine-grained clock synchronization, 2018.
2007, pp. 171–180.

Thank you