Time Capability in NETCONF

draft-mm-netconf-time-capability-00

http://tools.ietf.org/html/draft-mm-netconf-time-capability

Tal Mizrahi, Yoram Moses

Technion – Israel Institute of Technology

NETCONF, IETF Meeting, Berlin, July 2013

Overview

- This draft defines the **time** capability.
- Allows time-triggered configuration updates.
- Client can attach scheduled time of execution to each RPC.
- Server can attach timestamp to RPC reply.

- This draft is part of a work-in-progress research of time-based configuration updates.
- A similar extension was proposed to the Open Networking Foundation (ONF) in the context of OpenFlow.

Example 1: Switch to Candidate Datastore

- Apply the candidate datastore to all devices.
- Send n commit messages.

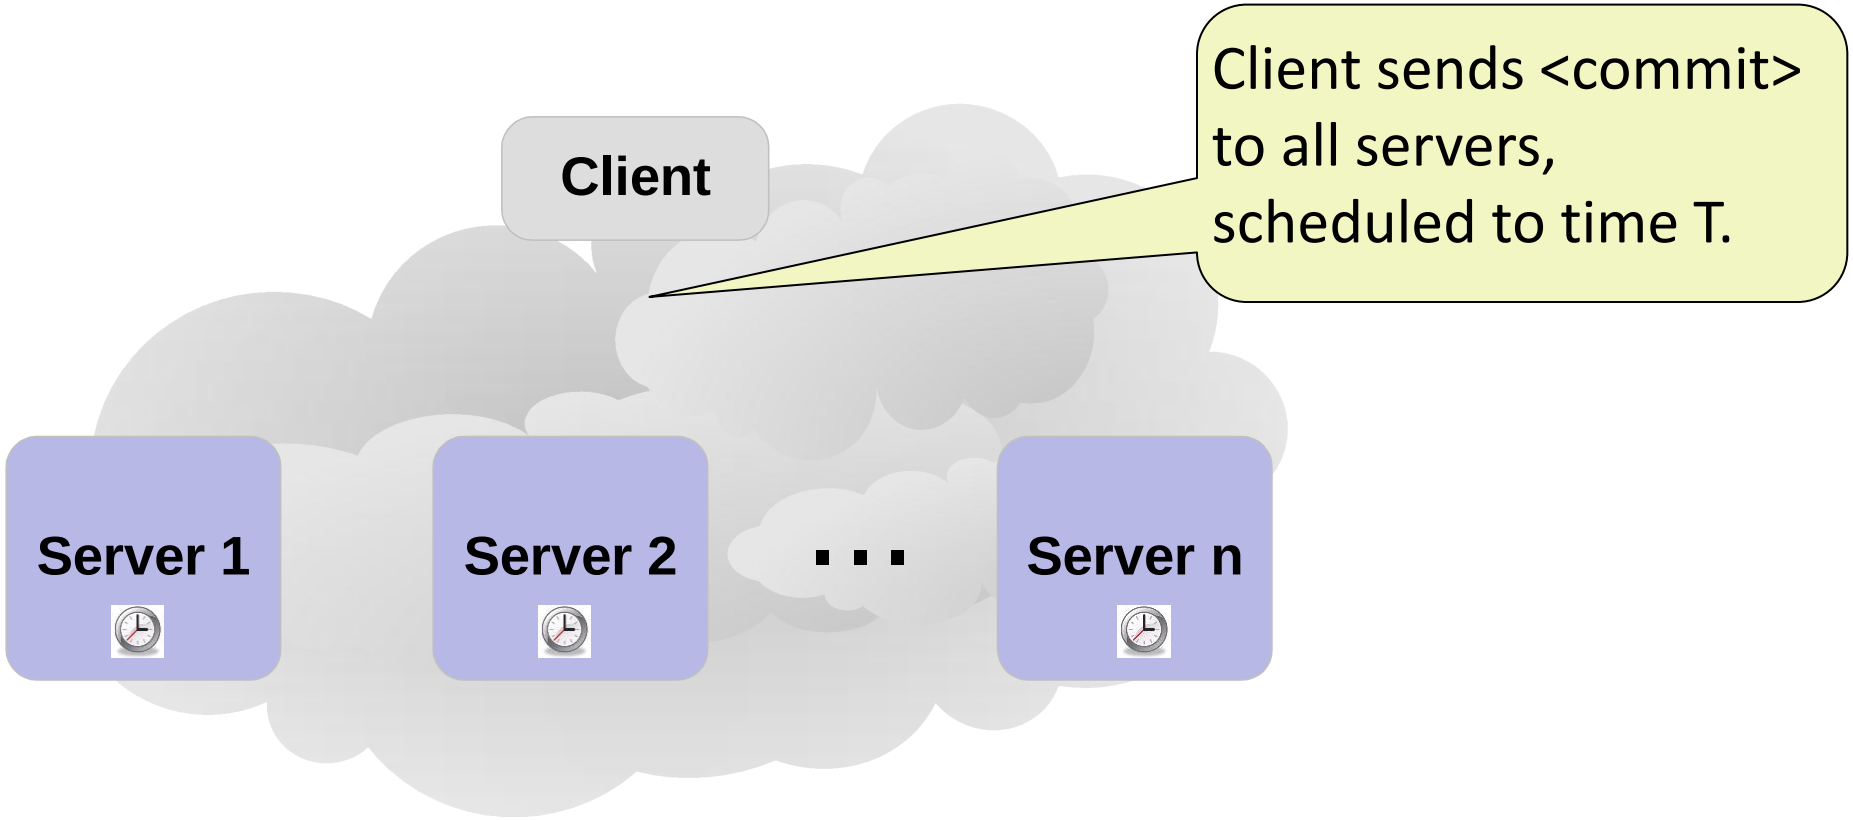

Example 2: Reconfigure Port Rate

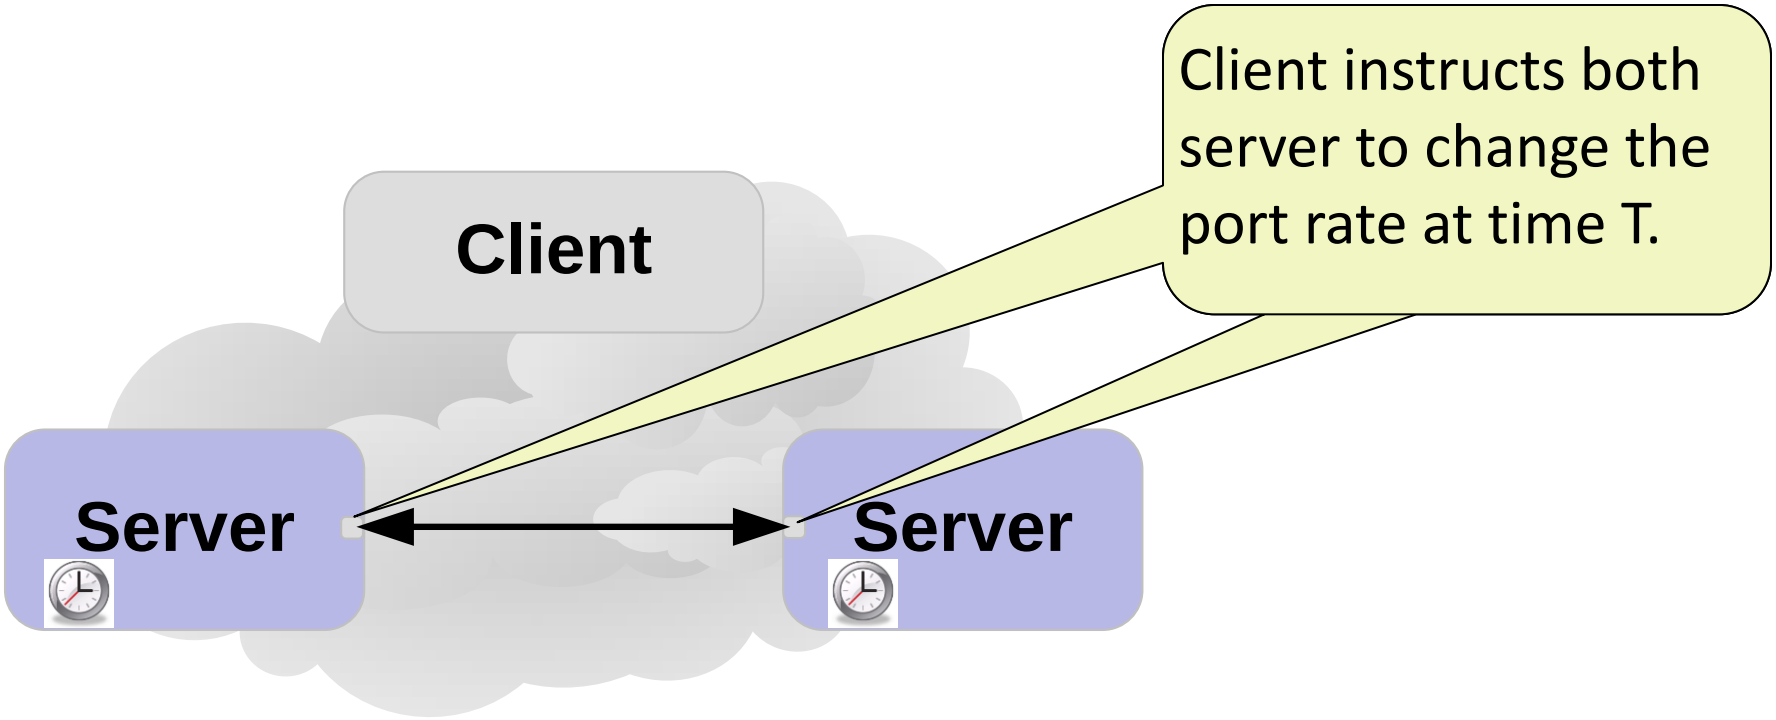

Example 3: Routing Change

- I2RS: rapid and dynamic routing changes
  (e.g., http://tools.ietf.org/html/draft-atlas-i2rs-problem-statement-00).
- This example: update the route to AS1.

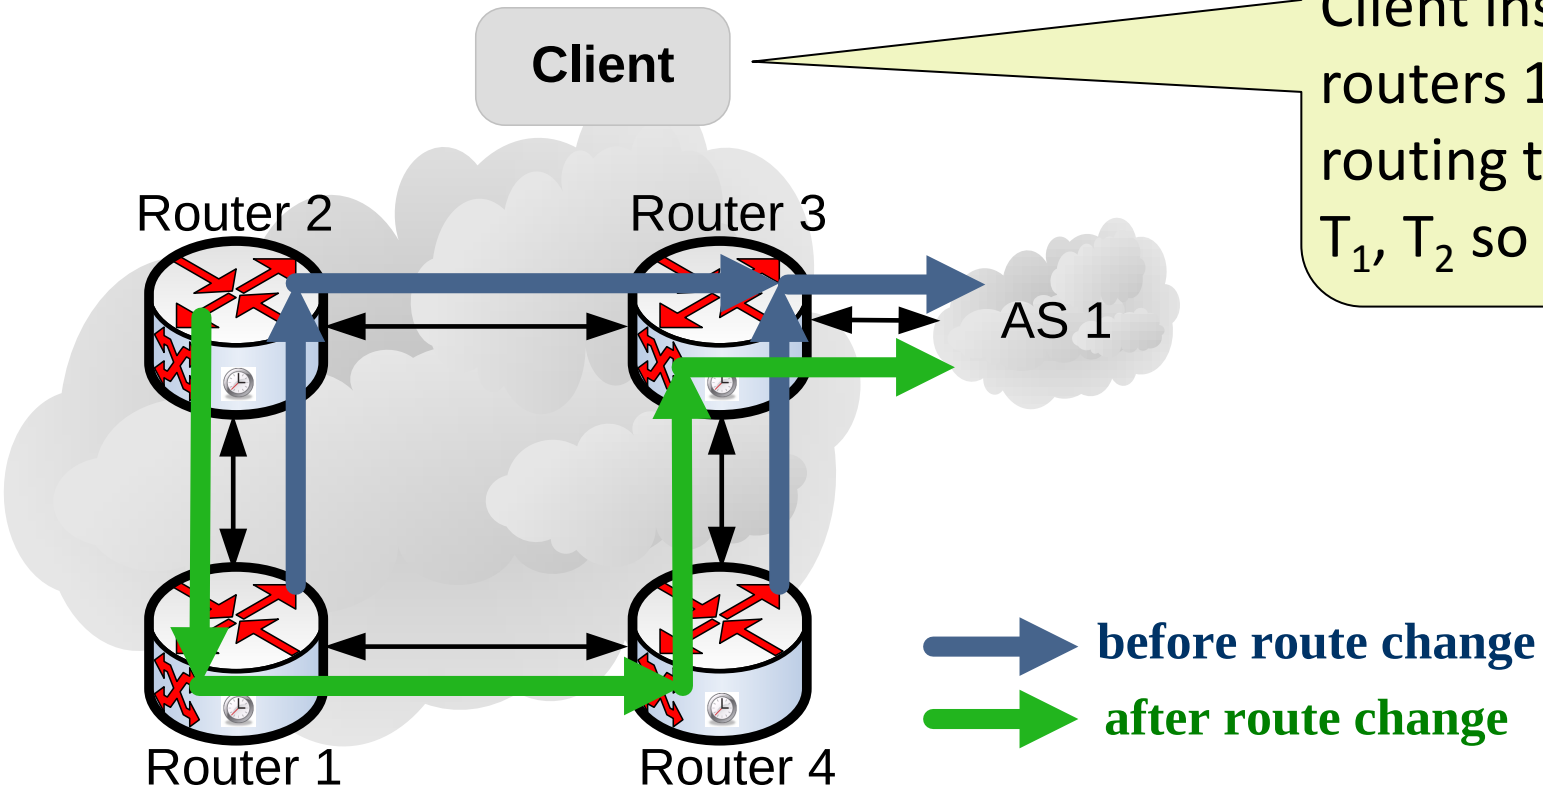

Scheduled Operations

- RPC may include <scheduled-time> element.

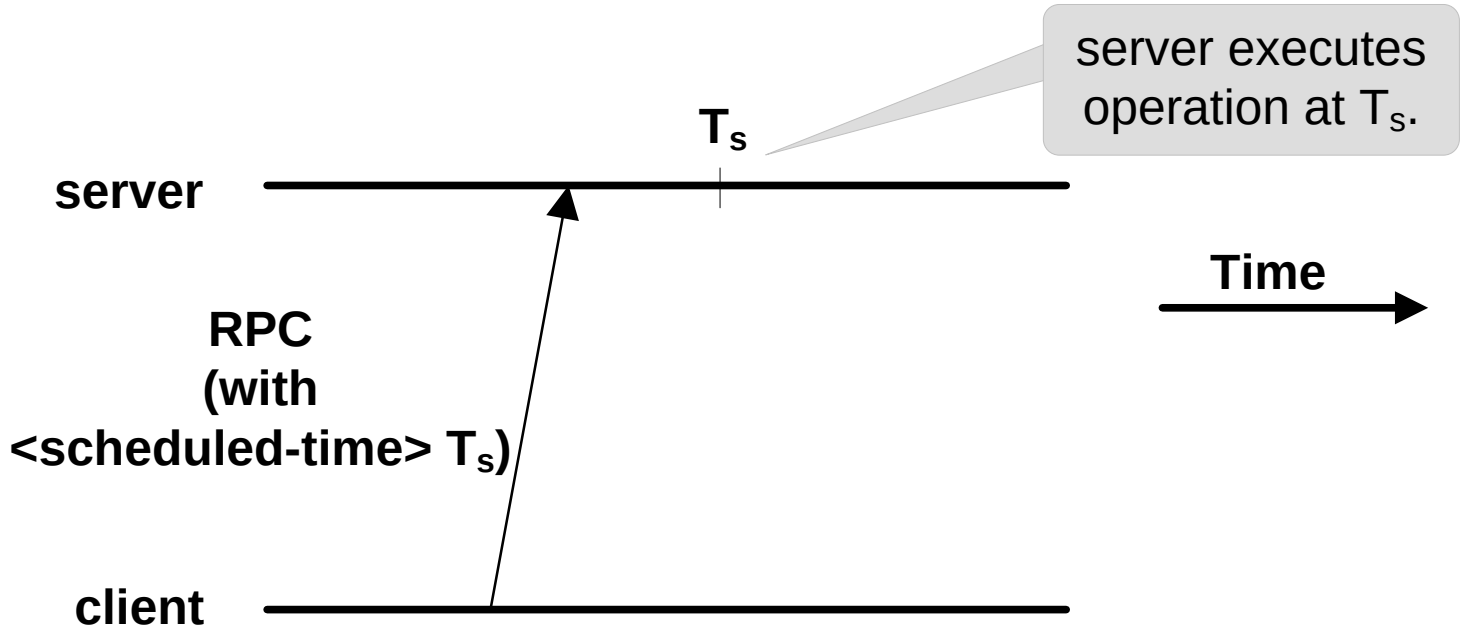

Reporting the Execution time

- RPC may include <get-time> element.
  ➔ RPC reply includes <execution-time>.

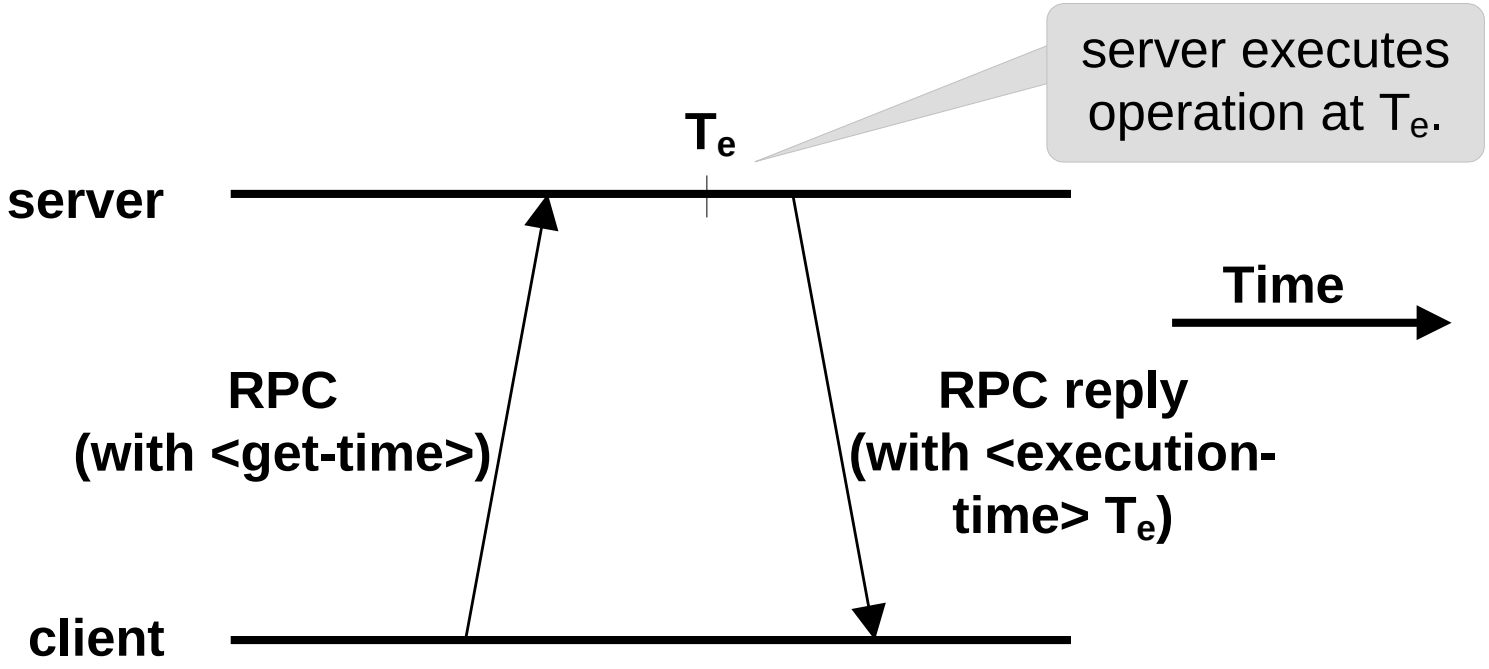

Scheduling and Reporting

- RPC may include both <get-time> and <scheduled-time>.
- Client receives feedback about whether the operation was executed as scheduled.

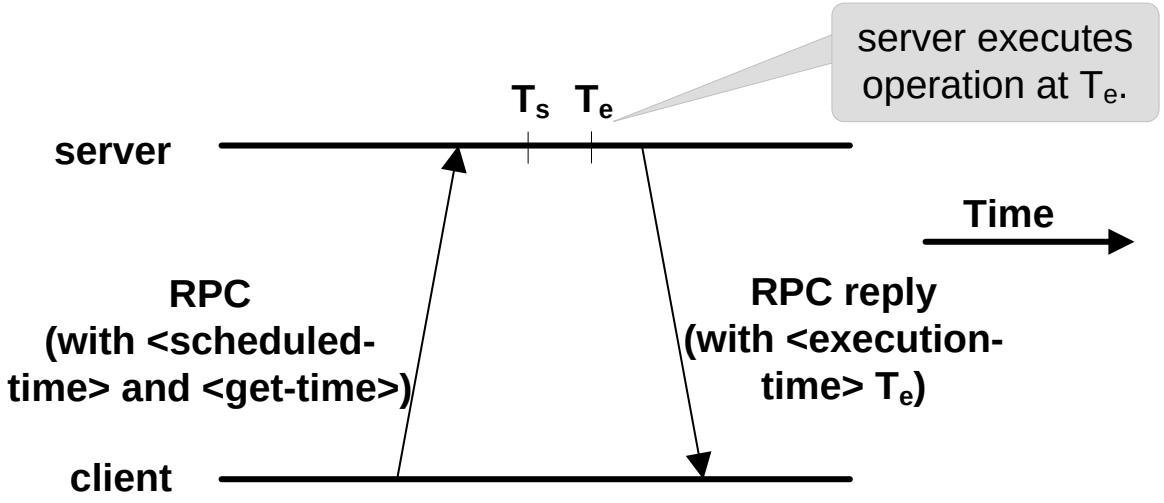

Summary

- Time-based updates can be used for:
  - Reducing transition period.
  - Physical layer updates.
  - Time-based sequence of ordered updates.

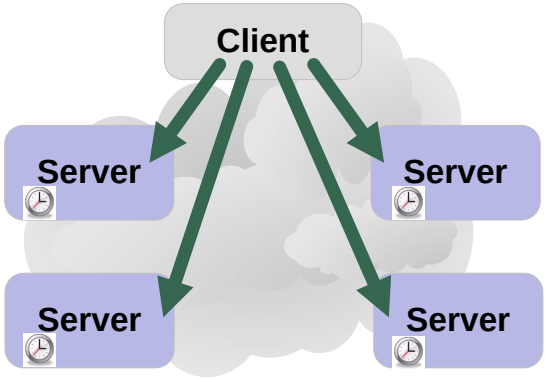

Issues Raised on the Mailing List

Question/comment	Response
Only relevant to <commit> or <edit-config> on :writable-running.	May also be relevant to <get>, <lock>, <unlock>, and future operations.
YANG date-and-time is better time parameter than 'seconds since 1970'.	IETF uses various time formats. We chose PTP time format.
Accuracy	Provision for a higher resolution than is currently needed
draft-kwatsen-conditional-enablement	draft-mm-netconf-time-capability allows: -Timed update for all RPCs. -Coordinated commit / get. -Sub-second resolution.

Issues Raised on the Mailing List (2)

Question/comment	Response
How would you see this working when supporting the configuration of a set of network elements in a robust and transaction-oriented way, where the operation should complete on all devices or be fully reversed?	To be addressed in next draft
A better solution would be an immediate <rpc-reply> (scheduled OK) and the execution results sent in a <notification>.	
Does access control get checked twice? Clients should not be able to schedule operations they are not permitted to execute.	
What if a session is lost or closed before its scheduled operation is started?	
What if the server reboots while operations are pending?	
How does a client cancel an operation?	
Can client A cancel operations for client B, assuming client A is allowed to invoke <kill-session>?	

History and Next Steps

- Draft 00 – July 2013.
- Next step: consider adding this topic to the WG charter.

THANKS !

Further Reading

- Mizrahi, T., Moses, Y., "Time-based Updates in Software Defined Networks", the second workshop on hot topics in software defined networks (HotSDN), to appear, 2013.
  http://tx.technion.ac.il/~dew/TimeSDN.pdf
- Mizrahi, T., Moses, Y., "Time-based Updates in OpenFlow: A Proposed Extension to the OpenFlow Protocol", Technion - Israel Institute of Technology, technical report, TR-1301, 2013.
  http://tx.technion.ac.il/~dew/OFTimeTR.pdf

BACKUP SLIDES

Example 4: Queue <max-rate> Reconfiguration

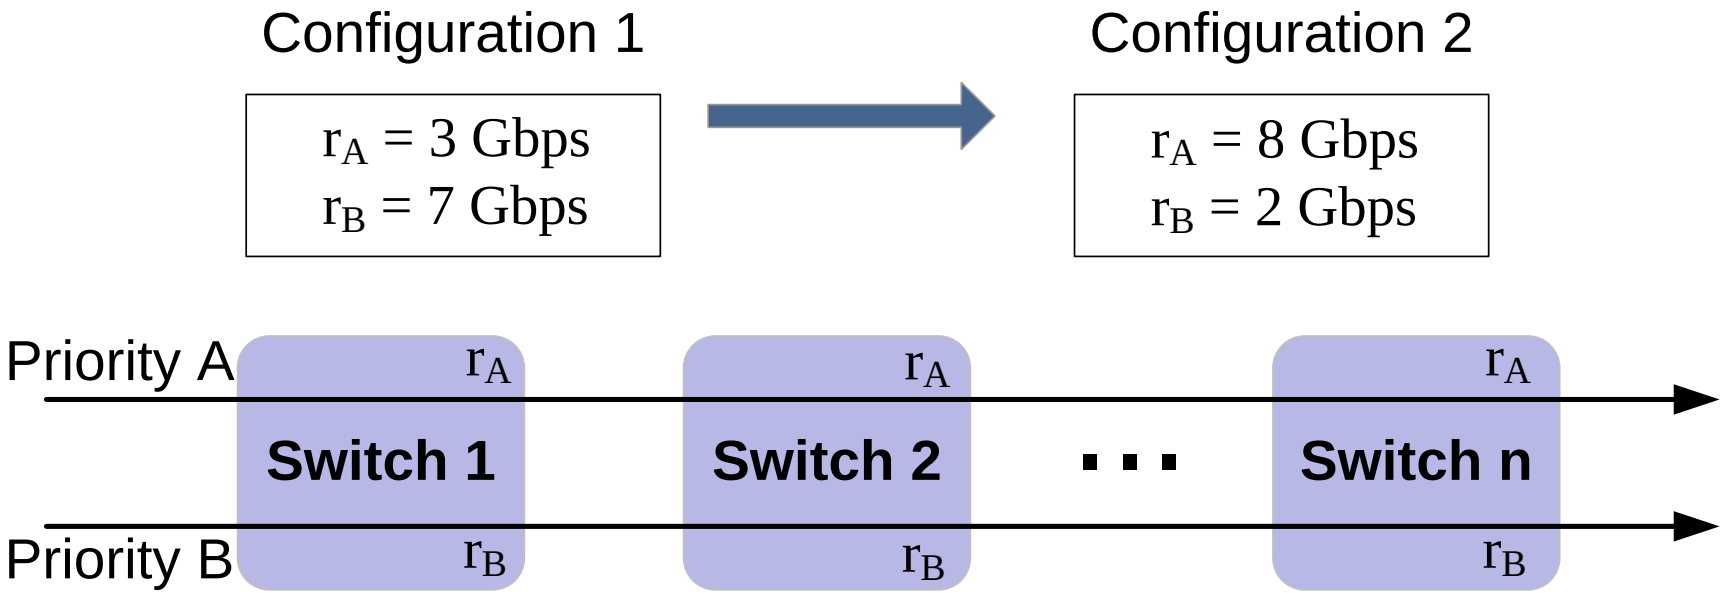

Using time: configuration point sends an updated <max-rate> to the n switches, scheduled to time T.

Example 5: Spanning Tree Reconfiguration

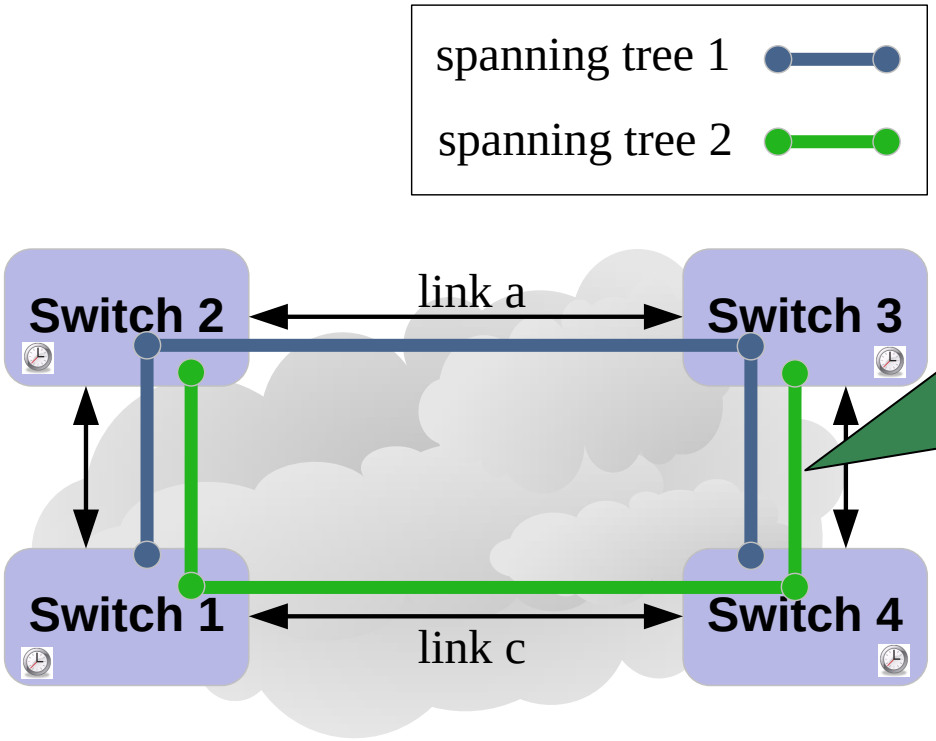

Using time:

- **Switches are configured to start using ST 2 at time T.**
- **Loops / loss of connectivity for a short transition period.**