

Root Cause Analysis:

Delayed watches and user updates

The following is a detailed accounting of the service outage that Rally users experienced on March 19-27th, 2019.

Root Cause Analysis Summary

Event Date	3/19/2020
Event Start	3/19/2020 18:00:00 MDT
Time Detected	3/20/2020 Approximately 09:00:00 MDT
Time Resolved	3/27/2020 Approximately 22:00 MDT
Event End Time	3/27/2020 Approximately 22:00 MDT
Root Cause	A subscription merge placed an unusually large number of change messages in the queue and consumers could not process the influx of messages in a timely manner.
Customer Impact	Delayed Watches notification processing for all subscriptions in certain partitions as well as login issues

Future Preventative Measures

Actions that should be taken to prevent this Event in the future.

Actions	Description
Improve documentation	Improve ops documentation around Watches notifications and messaging bus to include more information
Duplicate OCMs	Stop generating duplicate user object change messages (OCMs)
Lag monitoring	Spike on better monitoring around lag of messaging bus consumer topics as they related to micro-services
Monitor DB connection contention	DB connection contention monitoring for micro-services. Standardize this for use for others in the future.
Config options for Watches service connections	Implement config options for number of connections to Watches service so we don't have to change code to do it every time
Too many OCMs	Spike on subscription merge process for generating OCMs (just generate OCMs for new incoming objects)