

	ID	Date	Duration (hours)	Services Impacted	Triggers	Sustaining Effect	Mitigation
Google	GGL1 [8]	03/12/19	4.17	Gmail, Photos, Drive, Cloud Storage, various other GCP services	<ul style="list-style-type: none"> <li>• load spike</li> <li>• config change</li> </ul>	<ul style="list-style-type: none"> <li>• cascading overload</li> </ul>	<ul style="list-style-type: none"> <li>• load shedding</li> <li>• stop config deploy</li> </ul>
	GGL2 [9]	10/31/19	21.5	multiple components of GCE	<ul style="list-style-type: none"> <li>• software bug</li> </ul>	<ul style="list-style-type: none"> <li>• retry</li> </ul>	<ul style="list-style-type: none"> <li>• load shedding</li> <li>• reboot</li> <li>• capacity increase</li> </ul>
	GGL3 [10]	04/08/20	3.2	Google BigQuery, Cloud IAM 3% of Cloud SQL HA	<ul style="list-style-type: none"> <li>• config change</li> <li>• software bug</li> </ul>	<ul style="list-style-type: none"> <li>• retry</li> </ul>	<ul style="list-style-type: none"> <li>• config rollback</li> <li>• policy change</li> </ul>
	GGL4 [7]	04/30/13	1.5	Google API infrastructure	<ul style="list-style-type: none"> <li>• config change</li> <li>• latent software bug</li> </ul>	<ul style="list-style-type: none"> <li>• traffic queue growth</li> <li>• reboots</li> </ul>	<ul style="list-style-type: none"> <li>• config rollback</li> <li>• server reboot</li> </ul>
AWS	AWS1 [13]	04/21/11	66.7	Amazon EC2, Amazon RDS	<ul style="list-style-type: none"> <li>• network config change</li> </ul>	<ul style="list-style-type: none"> <li>• retry</li> </ul>	<ul style="list-style-type: none"> <li>• config rollback</li> <li>• policy change</li> <li>• load shedding</li> <li>• capacity increase</li> </ul>
	AWS2 [14]	06/13/14	4.23	Amazon SimpleDB	<ul style="list-style-type: none"> <li>• power loss</li> </ul>	<ul style="list-style-type: none"> <li>• retry</li> </ul>	<ul style="list-style-type: none"> <li>• load shedding</li> <li>• server restart</li> </ul>
	AWS3 [15]	09/20/15	4.55	AWS SQS, EC2 Autoscaling, CloudWatch, AWS Console	<ul style="list-style-type: none"> <li>• load spike</li> <li>• network disruption</li> </ul>	<ul style="list-style-type: none"> <li>• retry</li> <li>• cascading server demotion</li> </ul>	<ul style="list-style-type: none"> <li>• load shedding – pause metadata ops</li> <li>• capacity increase</li> </ul>
	AWS4 [16]	12/07/21	9.3	AWS DynamoDB, EC2, Fargate, RDS, EMR, Workspaces, AWS Console, Authorization services, internal DNS	<ul style="list-style-type: none"> <li>• latent software bug triggered by scale-up led to load spike</li> </ul>	<ul style="list-style-type: none"> <li>• retry</li> </ul>	<ul style="list-style-type: none"> <li>• load rebalancing</li> <li>• load shedding</li> </ul>
Azure	AZR1 [2]	07/01/20	2.65	Azure SQL DB & SQL Data Warehouse, Azure Database for MySQL/PostgreSQL/MariaDB	<ul style="list-style-type: none"> <li>• unspecified load imbalance trigger</li> <li>• latent config bug</li> </ul>	<ul style="list-style-type: none"> <li>• cascading overload</li> </ul>	<ul style="list-style-type: none"> <li>• service restart</li> </ul>
	AZR2 [2]	04/01/21	1.15	Azure DNS	<ul style="list-style-type: none"> <li>• software bug leading to cache degradation</li> </ul>	<ul style="list-style-type: none"> <li>• retry</li> </ul>	<ul style="list-style-type: none"> <li>• unknown automation</li> <li>• capacity increase</li> </ul>
	AZR3 [2]	06/14/21	13.25	Management operations of many Azure Services	<ul style="list-style-type: none"> <li>• latent software bug</li> <li>• load spike</li> </ul>	<ul style="list-style-type: none"> <li>• unspecified queue growth due to overload and timeouts</li> </ul>	<ul style="list-style-type: none"> <li>• load shedding</li> <li>• remove buggy software</li> <li>• capacity increase</li> </ul>
	AZR4 [2]	07/12/21	7.92	Windows Virtual Desktop, Azure Front Door, Azure CDN Standard	<ul style="list-style-type: none"> <li>• deployment of software bug</li> <li>• load spike</li> </ul>	<ul style="list-style-type: none"> <li>• retry</li> <li>• other unspecified</li> </ul>	<ul style="list-style-type: none"> <li>• load rebalancing</li> <li>• trigger hot fix</li> <li>• policy change</li> </ul>
Other	IBM1 [4]	06/11/21	73.53	Private DNS, HS Crypto Service, Cloudant DNS Services, Osaka, Cloudshell services	<ul style="list-style-type: none"> <li>• software bug</li> </ul>	<ul style="list-style-type: none"> <li>• retry</li> </ul>	<ul style="list-style-type: none"> <li>• load shedding</li> <li>• policy change</li> <li>• trigger hot fix</li> </ul>
	SPF1 [6]	04/13	NA	core app/service UI	<ul style="list-style-type: none"> <li>• load spike</li> <li>• policy failure</li> </ul>	<ul style="list-style-type: none"> <li>• retry</li> </ul>	<ul style="list-style-type: none"> <li>• load shedding</li> </ul>
	SPF2 [6]	06/04/13	8.33	core app/service UI	<ul style="list-style-type: none"> <li>• load spike due to unexpected service dependency</li> </ul>	<ul style="list-style-type: none"> <li>• retry</li> <li>• excessive logging in failure case</li> </ul>	<ul style="list-style-type: none"> <li>• trigger hot fix</li> <li>• load shedding</li> </ul>
	ELC1 [11]	04/02/19	6.67	Elasticsearch Service	<ul style="list-style-type: none"> <li>• unspecified maintenance</li> <li>• unspecified error</li> </ul>	<ul style="list-style-type: none"> <li>• load caused ZK churn causing more load</li> </ul>	<ul style="list-style-type: none"> <li>• restart</li> <li>• load shedding</li> </ul>
	WIK1 [17]	03/30/21	2.25	media upload, misc queued jobs	<ul style="list-style-type: none"> <li>• load spike</li> </ul>	<ul style="list-style-type: none"> <li>• unspecified causing queue growth</li> </ul>	<ul style="list-style-type: none"> <li>• load shedding</li> <li>• policy change</li> </ul>
	CCI1 [3]	07/07/15	18.33	Core product	<ul style="list-style-type: none"> <li>• load spike</li> </ul>	<ul style="list-style-type: none"> <li>• load increase due to contention</li> </ul>	<ul style="list-style-type: none"> <li>• load shedding</li> </ul>
	CAS1 [1]	07/27/17	NA	Partial database outage	<ul style="list-style-type: none"> <li>• rolling restart</li> </ul>	<ul style="list-style-type: none"> <li>• self-sustaining and increasing overload</li> </ul>	<ul style="list-style-type: none"> <li>• policy change</li> </ul>
	CAS2 [12]	2020	0.16	ably services	<ul style="list-style-type: none"> <li>• load spike of certain costly operations</li> </ul>	<ul style="list-style-type: none"> <li>• retry</li> </ul>	<ul style="list-style-type: none"> <li>• trigger removal – operated in stable state</li> </ul>
FB1 [5]	NA	NA	Facebook core services	<ul style="list-style-type: none"> <li>• load spike</li> </ul>	<ul style="list-style-type: none"> <li>• software bug</li> </ul>	<ul style="list-style-type: none"> <li>• hot fix</li> </ul>	

Table 1: Metastable failures from public sources. Azure and IBM do not provide a direct incident link. Gray highlight indicates a plausible metastable failure, although the incident description lacked some necessary details.

## References

- [1] Anonymous. Overload because of hint pressure + MVs. Apache Cassandra Issue Tracker: <https://issues.apache.org/jira/projects/CASSANDRA/issues/CASSANDRA-13810?filter=allopenissues>, 2017.
- [2] Microsoft Azure. Azure status history. <https://status.azure.com/en-us/status/history/>, 2021.
- [3] CircleCI. DB Performance Issue Incident Report for CircleCI. <https://circleci.statuspage.io/incidents/hr0mm9xmm3x6>, 2015.
- [4] IBM Cloud. Incident reports. <https://cloud.ibm.com/status/incident-reports>, 2021.
- [5] Facebook. Solving the Mystery of Link Imbalance: A Metastable Failure State at Scale). <https://engineering.fb.com/2014/11/14/production-engineering/solving-the-mystery-of-link-imbalance-a-metastable-failure-state-at-scale/>, 2020.
- [6] David Pobladori Garcia. Incident Management at Spotify. <https://engineering.atspotify.com/2013/06/04/incident-management-at-spotify/>, 2013.
- [7] Google. Google API infrastructure outage incident report. Google Developers blog: [https://developers.googleblog.com/2013/05/google-api-infrastructure-outage\\_3.html](https://developers.googleblog.com/2013/05/google-api-infrastructure-outage_3.html), 2013.
- [8] Google. Google App Engine Incident #19007. <https://status.cloud.google.com/incident/appengine/19007>, 2019.
- [9] Google. Google Compute Engine Incident #19008. <https://status.cloud.google.com/incident/compute/19008>, 2019.
- [10] Google. Google Cloud Infrastructure Components Incident #20005. <https://status.cloud.google.com/incident/zall/20005>, 2020.
- [11] Panagiotis Moustafellos and Ben Osborne. Elastic Cloud Incident Report: February 4, 2019. <https://www.elastic.co/blog/elastic-cloud-incident-report-february-4-2019>, 2019.
- [12] Byers Paddy. Cassandra counter columns: nice in theory, hazardous in practice. <https://ably.com/blog/cassandra-counter-columns-nice-in-theory-hazardous-in-practice>, 2021.
- [13] Amazon Web Services. Summary of the Amazon EC2 and Amazon RDS Service Disruption in the US East Region. <https://aws.amazon.com/message/65648/>, 2011.
- [14] Amazon Web Services. Summary of the Amazon SimpleDB Service Disruption. <https://aws.amazon.com/message/65649/>, 2014.
- [15] Amazon Web Services. Summary of the Amazon DynamoDB Service Disruption and Related Impacts in the US-East Region. <https://aws.amazon.com/message/5467D2/>, 2015.
- [16] Amazon Web Services. Summary of the AWS Service Event in the Northern Virginia (US-EAST-1) Region. <https://aws.amazon.com/message/12721/>, 2021.
- [17] Wikitech. Incident documentation/2021-03-30 Jobqueue overload. [https://wikitech.wikimedia.org/wiki/Incident\\_documentation/2021-03-30\\_Jobqueue\\_overload](https://wikitech.wikimedia.org/wiki/Incident_documentation/2021-03-30_Jobqueue_overload), 2021.