

**Outline**


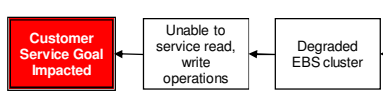
<b>What</b>	Problem(s)	Loss of service
<b>When</b>	Date	April 21, 2011
	Time	12:47 AM PDT
	Different, unusual, unique	Traffic shift executed incorrectly
<b>Where</b>	State, city	US East Region Availability Zone
	Facility, site	Elastic Compute Cloud (EC2)
	Unit, area, equipment	Elastic Block Store (EBS)
	Task being performed	Upgrade capacity of primary network

**Impact to the Goals**

<b>Safety</b>	N/A
<b>Environmental</b>	N/A
<b>Cust. Service</b>	Unable to service read, write operations
<b>Production-Schedule</b>	?
<b>Property, Equip, Mtls</b>	N/A
<b>Labor, Time</b>	Time, effort to fix problem

Frequency	Rare	This incident ?
		Annualized Cost ?

**Cause Map**  
Detail Level



**Solutions**

No.	Action Item	Cause
1	Reroute data to primary network	Data shifted from primary to secondary network
2	Increase capacity	High volume of nodes searching for replicas
3	Modify retry logic	Long time-out period
4	Identify, repair source of race condition	?

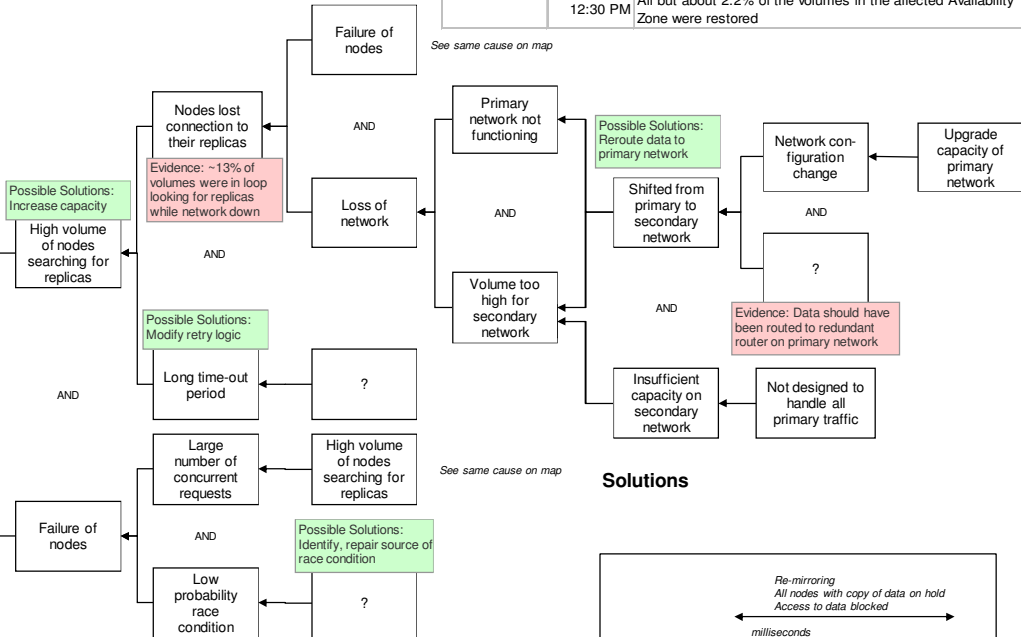
**Loss of Network Cloud Compute Service US East Region April 21, 2001**



Copyright ThinkReliability 2011

**Timeline**

Date	Time	Description
April 21, 2011	12:47 AM	Configuration change to upgrade the capacity of the primary network
	2:40 AM	Team deployed a change that disabled all new Create Volume requests in the affected Availability Zone
	8:20 AM	Team began disabling all communication between the degraded EBS cluster in the affected Availability Zone and the EBS control plane
	11:30 AM	Team developed a way to prevent EBS servers in the degraded EBS cluster from futilely contacting other servers
	12:04 PM	Outage was contained to the one affected Availability Zone and the degraded EBS cluster was stabilized
April 22, 2011	2:00 AM	Team successfully started adding significant amounts of new capacity and working through the replication backlog
	12:30 PM	All but about 2.2% of the volumes in the affected Availability Zone were restored



**Process Map**

