# RedShield Cloud
# Managed Application Security Service


## INCIDENT REPORT

## CASE # 49951

# Purpose of this Document

This document contains analysis and follow up relating to a severity 1 incident tracked in multiple RedShield support cases.

# Version Control

| Version: | Created by: | Date: | Description: |
|---|---|---|---|
| 0.1 | Daya Rajaratnam | 29/7/2020 | Initial Draft |
| 1.0 | Sam Pickles | 29/7/2020 | Review |

# Confidentiality

This document must be treated as Commercial in Confidence, and distribution is restricted to RedShield and its direct customers.

# Disclaimers as to the Information

The information contained in this document is for general information purposes only. RedShield assumes no responsibility for the information provided by its third-parties, technologies and tools used to arrive at the information presented.

# Review and Approval

| Name: | Details: |
|---|---|
| Sam Pickles | CTO, RedShield Security Ltd<br>79 Boulcott Street, 12th Floor, Wellington, NZ 6011<br>sam@redshield.co |

# Contact

All enquiries relating to this document should be referred to:

| Name: | Details: |
|---|---|
| Daya Rajaratnam | Manager Infrastructure, RedShield Security Ltd<br>79 Boulcott Street, 12th Floor, Wellington, NZ 6011<br>daya@redshield.co |

# Table of Contents

# 1. Management Summary

A partial service outage was caused by an F5 software bug, which caused some customer sites to fail between 14:00 to 18:00 UTC. RedShield's monitoring systems failed to immediately detect the issue, and automatic failover did not trigger.

| | |
|---|---|
| **Incident Ticket:** | 49932, 49939, 49944, 49951, 49950 |
| **Severity:** | Critical |
| **Services Impacted:** | Web services via Melbourne Datacenter |
| **Service Impact Period:** | July 28th 2020 14:00 to 18:00Hrs UTC |
| **Application Impacted:** | Traffic services via Melbourne Datacenter cluster 1 |
| **Incident Cause:** | Database ran out of allocated space due to extensive logging |
| **Incident Resolution** | Reboot, failover and reduce the amount of logging |
| **Outcomes** | Cluster is operational with an adequate allocation of disk space. An additional monitoring system has been brought into service. |

# 2. Root Cause Analysis

The active F5 cluster node at the Melbourne datacenter stopped processing traffic intermittently due to excessive logging, which caused disk write error messages and database allocated disk space overrun. Both units of the cluster went active; causing the cluster to serve some traffic periodically, and to fail to pass some traffic at other times, for the duration of the incident until manual intervention.

Disk space utilisation overall was not excessively high, which meant that SNMP monitoring did not trigger an alarm. The F5 device remained active, and responded to ICMP ping and TCP monitoring probes, which meant that cluster status initially appeared normal. Global datacenter failover also did not trigger as the failed F5 device was reporting normal status.

# 3. Resolution and Recovery

The failed cluster node needed to be manually rebooted to migrate the active sessions off of that defective device, by rebooting the device. This caused the cluster to fail-over the remaining sessions

to the standby cluster node and to restore the sessions and prevent further flapping.

# 4. Corrective and Preventive Measures

- The F5 support team is engaged and looking into why there were high amounts of log messages and looking at ways to limit the number of log files, so the disk space usage is limited.
- RedShield have implemented enhanced monitoring and escalation to detect this type of failure, in addition to the above preventive measures taken.
- Specifically, RedShield's Infrastructure team has reengineered our Content Route monitoring system to provide end-to-end service checking using HTTPS request probes, for every customer application globally, which continuously determine availability based on server responses through both primary and secondary RedShield datacenters. Alerts are correlated continuously to determine when a shield cluster is degraded and automatic failover has not occurred, and trigger an immediate service team response through our escalation system (ZenDesk/PagerDuty). Successful testing of this monitoring and escalation system was completed in production on Wednesday at 9PM.
- This enhanced monitoring and escalation system is designed to quickly detect the type of issue we have experienced here, in any RedShield cluster globally; in which the shielding cluster (F5 Networks Big IP) is degraded for a significant number of applications, but does not automatically fail over the cluster locally, or trigger a datacenter failover globally via dynamic DNS; and continues to report status green via SNMP, ICMP and other existing monitor probe types.