Banner Outage

Problem Solving Analysis

Event Occurrence: December 22, 2017

Background

At approximately 11am on Friday, 12/22/17, production Banner and reporting instances experienced significant performance degradation. Performance degraded to such a state that Database Administrators opted to shutdown Banner in order to protect data and avoid potentially significant data corruption or loss.

Break Down of the Problem

The enterprise disk array serving production Banner and RPTP had two failed disks. Due to a parts shortage reported by the vendor, technical staff, without communicating to nor consulting with acting management, weighed risks and opted to replace the failed drives with spares from another non-production array.

Unbeknownst to the technical staff at the time, one of the spare replacement drives selected for this procedure was bad. As the array attempted to add the drives into its grouping, performance of production service databases, namely Banner PROD and RPTP, were negatively impacted. Services degraded to such a point that Database Administrators opted to shutdown Banner PROD to avoid data corruption.

Target State / Goal

1. The arrays under contract for support never reach a state of more than a single failed hard drive at one time.
2. OIT technical staff do not swap out hard drives in arrays for which a support contract exists in which the vendor is responsible for doing so.
3. “Old” or “used” parts are never selected for use in production enterprise systems.
4. Management is consulted with before significant decisions are made that could have large business impact.

Root Cause Analysis

Two hard drives had gone bad on EVA-03, which serves production Banner (PROD & RPTP). The drives were noted to be in a degraded state by TS staff on Wednesday, 12/20/17. HPE support was contacted and a case opened at that time. HPE Support advised there was a slight delay in obtaining the drives; however, they were able to have them on-hand in Anchorage on Thursday, 12/21/17. Due to the vendor calling and leaving voicemail directly on a technician’s desk phone (who was out sick that day), rather than contacting the NOC, the vendor was unable to contact someone to arrange for both physical access to the premises and to coordinate the software configuration changes that are necessary when replacing drives. As such, they were not able to travel to Fairbanks and conduct the work on Friday, 12/22/17. Because of this and believing that having more than two disks bad in the array at single time could lead to data corruption, TS staff opted to replace the two bad disks with spares from another non-production array between 10:30-10:45 that morning - a procedure that had been utilized successfully in the past. About 15 minutes later, the Service Desk began to receive reports of Banner slowdowns, with services continually degrading over time.

TS staff contacted HPE support for assistance in troubleshooting the degraded performance, opening case #5325720440. Due to the array degradation, it was difficult for the vendor to obtain the necessary logs. At 12:50 that afternoon, services had degraded to the point that production Banner was unusable. Database Administrators made the call to gracefully shutdown Banner PROD and RPTP while they still could, in order to avoid potential data corruption or loss. The vendor was eventually able to pinpoint the time the issues started, which correlated with the old spare drives being installed. With guidance from HPE’s Storage experts, technical staff removed the old disks between 2-2:30pm, restoring the array. After vendor validation, production Banner was made available at approximately 3:40pm.

During troubleshooting, TS was advised by the HPE Storage Engineer that there was no danger of data loss because we could have up to four failed drives at a single time. With this new information, the disk replacement was scheduled to after the holiday closure, as the support contract in place is for 8x5xNBD, and not weekend or holiday response/support. TS and HPE staff coordinated Wednesday morning, 1/3/18 and the two failed disks were replaced by the vendor on Thursday, 1/4/18 with no service degradation nor interruption.

Develop Countermeasures

1. Emphasize with staff to communicate with management during a service impacting event.
2. Validate with vendor proper contact procedures for on-site access and execution of work.
3. Be more proactive in not allowing multiple drives to be in a failed state at one time.
4. Do not attempt to use components of unknown state in enterprise systems/equipment.