Reliability Analysis of SAN Topologies under Variable Workloads Using Hybrid Simulation Techniques

Jacek Nazdrowicz[1], Maja Tuszynska[2]
[1] Lodz University of Technology/Fujitsu, Lodz, Poland
[2] Cracow University of Technology, Cracow, Poland
jacek.nazdrowicz@p.lodz.pl

Extended Abstract

Storage Area Networks (SANs) are widely deployed with redundant architectures to ensure high availability, yet their actual performance under realistic operating conditions - characterized by workload variability and stochastic component failures - remains insufficiently captured by classical availability-centric evaluation methods. This work applies a hybrid reliability–performance simulation framework to quantify how redundancy influences queue dynamics and delay tail behavior, which directly determine user-visible performance.

The analysis focuses on comparative results obtained for single-fabric and dual-fabric SAN topologies under multiple workload scenarios, including burst-dominated and stress-mix regimes combining high traffic variability with frequent failures. Instead of relying on average throughput or steady-state availability, the evaluation emphasizes percentile-based backlog and distribution-based delay metrics, which are more sensitive to transient overload phenomena.

A central result of this study is illustrated in Fig. 1, which shows the time evolution of the 95th-percentile queue length for both SAN architectures. Dual-fabric configurations consistently maintain lower percentile backlog levels than single-fabric systems, confirming that redundancy mitigates congestion by preserving alternative paths during failure events. However, Fig. 2 also reveals that under stress-mix workloads the percentile queue length continues to grow over time even when connectivity is preserved. This demonstrates that sustained workload variability and temporary capacity reduction can drive long-term backlog accumulation independently of availability loss. The result provides direct evidence that high availability does not guarantee congestion-free operation.

The persistence of congestion effects is further highlighted by the delay tail behavior shown in Fig. 2, which presents complementary cumulative distribution functions (CCDFs) of delay. For stress-mix and failure-heavy scenarios, both architectures exhibit heavy-tailed delay distributions, indicating a non-negligible probability of extreme backlog-induced delays. Dual-fabric SANs shift the CCDF downward relative to single-fabric systems, reducing the probability of severe delay events, yet the heavy-tail structure remains pronounced. This indicates that redundancy improves delay characteristics quantitatively but does not alter the fundamental tail behavior when workload variability exceeds effective service capacity.

Taken together, the results in Fig. 1 and Fig. 2 expose a critical limitation of availability-focused SAN evaluation. While dual-fabric architectures significantly enhance connectivity and reduce average congestion, they cannot fully suppress backlog growth and tail latency under highly dynamic workloads. Percentile queue lengths and delay CCDFs therefore emerge as essential complementary metrics to availability when assessing SAN reliability and performance under realistic operating conditions.

The presented findings demonstrate the practical value of hybrid reliability–performance simulation: only by jointly analyzing failures, workload dynamics, and queue evolution can designers accurately quantify the benefits and limits of redundancy in enterprise-class SAN deployment.

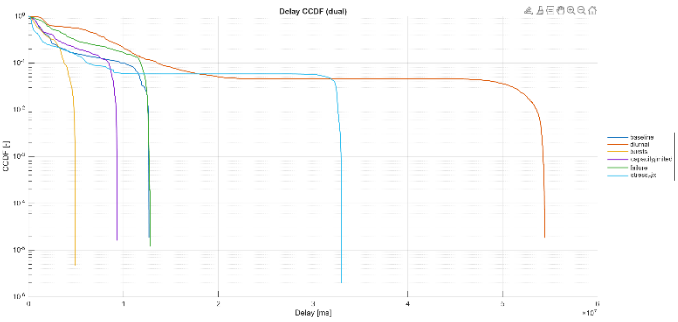

Fig. 1. Delay CCDF (dual-fabric). Complementary cumulative distribution function (CCDF) of request delay for the dual-fabric SAN, emphasizing tail behavior and backlog-induced delay under stress conditions.

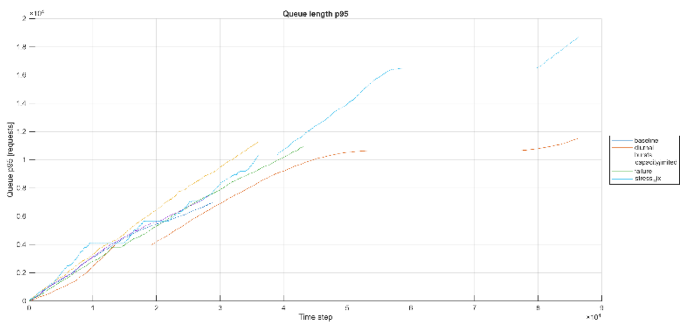

Fig. 2. Queue length p95 over time.