

## **Serverless Network File Systems – SOSP 1995**

1. What were the goals of xFS? How do these goals compare to Porcupine?
2. What assumptions did xFS make about their computing environment?
3. What assumptions did xFS make about workloads? What types of workloads do you think would run well on xFS?
4. xFS borrows heavily from RAIDs, LFS, and Zebra. Why is LFS a good match for RAIDs? How did Zebra combine LFS and RAIDs? Explain Figure 1. How does xFS extend beyond Zebra?
5. What is the primary design philosophy of xFS?
6. In a typical centralized system, the central server has four main tasks: storing all blocks on disk, managing disk location metadata to indicate where blocks are stored, caching blocks, and managing cache consistency. How are each of these four tasks distributed in xFS? Explain the operations shown in Figure 3.
7. Does the location of the manager for a particular file matter within the cluster? Where does the xFS design worry about locality and where does it not? Do you think this is the right choice?
8. How well do you think xFS can handle failures? What work must it do when a node fails? How well can xFS handle newly added nodes? How well can xFS handle performance heterogeneity?
9. What components of xFS were not implemented? Why not? Is there a lesson from this?
10. How did the designers think xFS would be used? How do you think xFS could be best used today?
11. Does the paper evaluate how well xFS meets its original goals? What experiments would you like to see to demonstrate this?
12. In their performance evaluations shown in Figures 9, 10, and 11, what are the workloads? Given their hardware, what is the best performance they could have achieved? What do you think of their comparison to NFS?
13. xFS distributes the task of the manager everywhere in the cluster. Do the experimental results shown in Figure 13 validate this decision?
14. What do you think is the contribution of xFS?