CHAPTER 1

Introduction

We live in a complex and uncertain world. Need we say more? However, we can say quite a bit about some aspects of randomness that govern behavior of systems—in particular, failure events. How can we predict failures? When will they occur? How will the system we are designing react to unexpected failures? Our task is to help identify possible failure modes, predict failure frequencies and system behavior when failures occur, and prevent the failures from occurring in the future. Determining how to model failures and build the model that represents our system can be a daunting task. If our model becomes too complex as we attempt to capture a variety of behaviors and failure modes, we risk making the model difficult to understand, difficult to maintain, and we may be modeling certain aspects of the system that provide only minimal useful information. On the other hand, if our model becomes too simple, we may leave out critical system behavior that dramatically reduces its effectiveness. A model of a real system or natural process represents only certain aspects of reality and cannot capture the complete behavior of the real physical system. A good model should reflect key aspects of the system we are analyzing when constrained to certain conditions. The information extracted from a good model can be applied to making the design of the system more robust and reliable.

No easy solutions exist for modeling uncertainty. We must make simplifying assumptions to make the solutions we obtain tractable. These assumptions and simplifications should be identified and documented since any model will be useful only for those constrained scenarios. Used outside of these constraints, the model will tend to degrade and provides us with less usable information. That being the case, what type of model is best suited for our project?

When designing a high availability system, we should carefully analyze the system for critical failure modes and attempt to prevent these failures by incorporating specific high availability features directly in the system architecture and design.

Designing High Availability Systems: Design for Six Sigma and Classical Reliability Techniques with Practical Real-Life Examples, First Edition. Zachary Taylor and Subramanyam Ranganathan.

However, from a practical standpoint, we know unexpected failures can and will occur at any time despite our best intentions. Given that, we add a layer of defense, known as fault management, that mitigates the impacts of a failure mode on the system functionality. Multiple failures and/or failure modes not previously identified may cause system performance degradation or complete system failure. It is important to characterize these failures and determine the expected overall availability of the system over its lifetime of operation.

Stochastic models are used to capture and constrain randomness inherent in all physical processes. The more we know about the underlying stochastic process, the better we will be able to model that process and constrain the impacts of the random failures on the system we are analyzing. For example, if we can assume that certain system components have constant failure rates, a wealth of tools and techniques are available to assist us in this analysis. This will allow us to design a system with a known confidence level of meeting our reliability and availability goals. Unfortunately, two major impediments stand in our way: (1) The failure rate of many of the components that comprise our system are not constant, that is, independent of time over the life of the system being built or analyzed, but rather these failure rates follow a more complicated trajectory over the lifetime of the system; and (2) exact component failure rates—especially for new hardware and software—are not known and cannot be exactly determined until after all built and deployed systems reach the end of their useful lives.

So, where do we start? What model can we use for high availability design and analysis? How useful will this model be? Where will it fail to correctly predict system behavior? Fortunately, many techniques have already been successfully used to model system behavior. In this book, we will cover several of the more useful and practical models. We will explore techniques that will address reliability concerns, identify their limitations and assumptions that are inherent in any model, and provide methods that in spite of the significant hurdles we face, will allow us to effectively design systems that meet high availability requirements.

Our first step in this seemingly unpredictable world of failures is to understand and characterize the nature of randomness itself. We will begin our journey by reviewing important concepts in probability. These concepts are the building blocks for understanding reliability engineering. Once we have a firm grasp on key probability concepts, we will be ready to explore a wide variety of classical reliability and Design for Six Sigma (DFSS) tools and models that will enable us to design and analyze high availability systems, as well as to predict the behavior of these systems.