

Courselanding

Hi and welcome to this ONTAP cluster training. My name is Peter and I will be instructor for this course. The prerequisite for this training could be that you either have some hands on experience in a NetApp cluster environment. All of that you have attended an ONTAP administration training. This is a somewhat advanced training that presupposes that you're familiar with the basic concepts. How has this training organized? The millennials will be a repeating pattern of theory followed by a lab in which we will break something and then fix it. Now what will you learn in this training? First, we'll start with an overview to refresh our knowledge of the ultra environment. We have to make sure that we're familiar with the common concepts and architecture before we can dig into any problem whatsoever. Then in the modules that follow, we will have repeating patterns in which we will describe a feature and possible problems that may occur. Then we will actually invoke a problem, think of a solution and implement that solution in order to fix whatever is broken. Module two is on the boot process. What happens when a node boots and how does it join the cluster? We will inspect the compact flash and also we will check the relevance of the file systems in there and how this plays a role in joining the cluster.

In this module we will also look at the most important options of the boot menu that have to do with configuration issues. Module three is on problems with the replicated databases. We will break and repair the RDB environment of a node. We'll also discuss and cause inconsistencies between WAFL and the volume location database and then we will fix that. Of course, module four is on the epsilon. Epsilon like the RDBMS is related to majority vote and Corum. It is only relevant if you run multiple HA pairs in your cluster. Then we will discuss problems with volumes. What happens if you lose the root volume of a node? What happens if you lose the root volume of a data storage virtual machine and what can go wrong with data volumes? Module six will be on networking issues. This module deals with the cluster interconnect as well as the data network failure of the groups and some more common things that can be problematic. Module seven, we'll deal with log files. Now remark about the lab environment. I will use the Netapp simulated to demonstrate everything. Sometimes I will use three nodes, something you will never see in real life, and sometimes I will use for notes. This means that you can do this training without the use of a physical environment. So let's start.