IBM Servers and Storage Demos
i5/OS Switch Disk Storage Architecture - i5/OS Independent Auxiliary Storage Pools

<0:00>
Welcome to Part 2 of Switch Disk Basics. In Part 1, we went over the fundamental storage architecture that led to the development of independent auxiliary storage pools. In Part 2, we will examine some of the advantages and uses of switchable independent auxiliary storage pools.

<0:22>
As we had reviewed before, an Independent Auxiliary Storage Pool is a storage pool that one can access from a given system. You can vary it on, vary it off, and when we make the Independent Auxiliary Storage Pool switchable, we give two systems the ability to vary on and attach that disk cluster in some way. So we can now switch between the two.

<0:51>
In this example, if System A is the Primary server, we go over to System A and we activate that IASP. In so doing, the disk is now varied on, attached, and active with System A. The application and the data associated with that application are running in this independent auxiliary storage pool. The advantage of making this switchable is that in the event we would like to do some system maintenance, an operating system upgrade, or we would like to apply fixes to an application or the operating system - and we cannot afford to take the system down for an extended period of time - what we can do is switch the disks from one system to another. We go to System A and Vary Off the independent auxiliary storage pool, and we make that switch disk cluster inactive.

<1:50>
Then we go over to System B and we vary it on. From that, the disk now becomes red because it belongs to System B. We can now run the same exact application and the same data from the disk using the processor and memory that are associated with System B. We can take as long as we need to, to do the kind of fixes or maintenance on System A. We did not lose any data in process because there was no data in flight. The data is committed and sitting on the independent auxiliary storage pool, so we always have a current copy of the data. In this way, we can switch back and forth between the two systems.

<2:36>
Most of the time these two systems have been depicted as separate servers connected via high-speed link. For many customers, managing two servers within the same data center with a disk cluster was a very viable option to respond to a shrinking maintenance window.

<2:56>
Other customers don't want to buy two separate servers and put them in the same data center. For those customers, instead of buying two servers, we would create two partitions on the same single server and now they can be switched between the partitions. This provides a tremendous amount of flexibility, because now we don't necessarily need to create two equally large servers in order to run the application. The second server can be in a fairly light weight partition relative to the size of the first one, because it doesn't need to do a lot of activity, until it becomes the active server. When it becomes the active server, we simply move the processing power over to that partition either in a scheduled or manual way, or we can take advantage of uncapped processor partitions on a Power 5 or later server. So there is a lot of flexibility, in terms of how we might configure the switch disk cluster.

<4:02>
Another issue that comes up for customers is the fact that system objects must reside in the system Auxiliary Storage Pool, which is non-switchable. System objects are things like system values, configurations, job descriptions, user names and profiles. Prior to Version 5 Release 4 of i5/OS, your options for managing these system objects was either do it manually or to invest in some software from a 3rd party who might manage it for you. With Version 5 Release 4 of i5/OS, IBM introduced something called 'Administrative Domain' or an 'Admin Domain'. The admin domain is part of cluster services, so when the customer purchases option 41 in Version 5 Release 4 or the High Availability solution in V6R1, one of the things they get is the admin domain capability. With the admin domain, I can define

different types of system objects - system values, user name, password. When I define these objects I call them managed resources. They will be automatically synchronized between the two systems. If I were to change a system value on either system A or System B, and that value is defined inside my admin domain, the change would be automatically synchronized across the two servers.

<5:39>
So admin domains are a very powerful way to manage those objects using synchronization technology as opposed to replication technology. That means whichever server made the last change, the change is the current state for that managed resource.

<5:57>
On this page you can see a list of managed resources available in V5R4. It includes system values, user profiles, job descriptions, network attributes, and many other things. Additional objects were added in V6R1 of i5/OS.

<6:15>
Another interesting aspect of this technology is that you can take a given object, like a user profile, and specify which parameters within the object are managed resources. Earlier in this discussion, I intentionally referred to user name and password, and not user profile. For example, administrators can determine that they only want the user name and password synchronized across servers; and not allow users to also synchronize their security levels between the two servers. There is a fairly good level of control over these objects that are supported in an admin domain.

END