LESSONS LEARNED AFTER IMPLEMENTATION AND MANAGEMENT OF HALF OF THE SNS DIAGNOSTICS PC-BASED INPUT OUTPUT CONTROLLERS (IOCS) [1]

J.D. Purcell, D. Murphy
ORNL, Oak Ridge Tennessee, USA

ABSTRACT

The Spallation Neutron Source (SNS) Project's Diagnostics Group has completed its implementation of half of the PC-based Network Attached Devices (NAD) [2]. The latest deployment included new NAD installations and configuration upgrades to existing NADs. This deployment included modifications to the base XP Embedded (XPE) OS and the addition of software to support patching, virus/worm protection, and database interaction. After deployment, these devices were used to support commissioning activities for the SNS linac. This paper describes the deployment commissioning experience and discusses plans for the future.

INTRODUCTION

The SNS project is a partnership involving six U.S. Department of Energy (DOE) national laboratories: Argonne, Brookhaven, Jefferson, Los Alamos, Lawrence Berkeley, and Oak Ridge. The project is described in more detail on the SNS Web site [3].

The SNS project has finished its commissioning of the Super Conducting LINAC. This accomplishment required the use of half of the diagnostic devices needed for the entire SNS project. With this completed work, the diagnostics group has now successfully installed and commissioned around 180 PC-based NADs.

Many factors were considered before deciding on the use of Microsoft Windows on a PC. From an operations and controls point of view, these IOCs had to interface and function with the other groups' systems. The operating system and platform had to be transparent to operations. The decisions to proceed with this plan may still be debated but we now have deployed systems that are all industrial rack-mounted PCs with a Microsoft Windows operating system. From the outside, they look and act like any other IOC and have proven to be successful. As we have purchased in stages, we have progressed from a 2.4GHz-processor on up. Each has 512 MB RAM with 40GB hard drives. These systems have been designed for use in a construction environment and allow for easy maintenance and proper cooling. The operating system is actually Microsoft's XPE with Service Pack 2. Many software and hardware tools are used to maintain the PCs, including Altiris Deployment Solution, ORACLE, Windows Server 2003, and simple remote power controllers.

During the installation of the first half of our devices, many changes have been made to the plans. The use of the PC and XPE has allowed us to stay on top of the issues. Our strategy is set up to handle these modifications easily, but it also creates some issues that were not planned initially. Even with these hiccups, the use of this strategy has given us a sound, stable technique for long-term use of the PCs.

DEPLOYMENT EXPLANATION

Why PC-Based?

We knew there were several factors to consider before choosing an approach for deployment. The original system architecture was going to be with the traditional Versa Module Eurocard (VME) bus and more specifically VME Extensions for Instrumentation (VXI). The VXI implementation would allow for multiple beam line devices to be controlled by one VME implementation. This, however, would be a single point of failure and moved us in a direction away from VXI and toward one input output controller (IOC) for each beam line device. The final design of these IOCs is one that tried to eliminate the brittleness of the VXI system and in turn employ a loose coupling between devices. The IOCs are now self-contained with their own pickups or sensors, and each IOC has its own resources for timing, data acquisition and data processing.

Using a rack-mounted PC forced us to decide on an operating system. When we started this project, we knew that over time our hardware would fail or that it would just have to be upgraded. This led us down a path that has actually given us answers to other challenges. Because our number one fear was the failure of our hard drives, we looked into an embedded OS and a plan was to deploy this OS on compact flash. This would eliminate the hard drive as a source of failure. Later we found that the hard drive was not the number one hardware failure but that we should actually be more concerned with the fans and our operating temperatures. With this realization, we moved back to the use of hard drives. We did not, however, eliminate the use of the embedded OS. At the time, our choice was between NT Embedded and embedded Linux. Three factors led us to start testing with NT Embedded: the fact that driver support for many of our PCI cards was not found in Linux, we wanted to stay with an off-the-shelf approach, and that MS Windows had progressed to a stable OS. We discovered that an embedded OS provided opportunities we had not really appreciated before.

Our use of off-the-shelf hardware and software eliminates a certain risk of future problems that might be encountered with the more traditional approach of VXI. With VXI, you are required to maintain certain specifications on your hardware. This means that it is the hardware that must be continually supported by the vendors, or you must prepare for hardware failure through the purchase of spare equipment. Our philosophy is that the specifications we adhere to are not specific to the hardware but more generic, a step above the PC, TCP/IP for example. Off-the-shelf materials will continue to follow these standards, and we will not have to prepare for materials that will not be supported by a vendor. Now, whatever the reason, systems can be upgraded, replaced, etc., and there are no adverse effects to operations.

XP Embedded

So now we deploy PCs with Windows XPE as the OS. XPE is Microsoft's XP Pro chopped up into a long list of components. And, as the XPE operating system is based on the same binaries as Windows XP Pro, we can build our OS out of only the components that we need. At first, we thought that we would maintain pretty small footprints for our systems, but we have found that we actually like to have a lot of the functionality of XP Pro. These machines are consistently used as development continues and support the addition of several different Windows-based applications. This actually has not affected us because we have returned to using standard hard drives and OS size is not a concern. Often, we can include components on the chance that we might need them. But XPE was also ultimately chosen as the OS platform because of its small footprint on disk and because of its ability to run with only specifically needed services. We have eliminated from our base OS image many of the potential network security concerns. We looked at XP Pro and realized that even some of our groups' administrative users might want to turn a service back on. With a service being eliminated completely, our security becomes that much stronger. Currently our embedded base OS stands at approximately 300 MB compared to approximately 1 GB for the basic XP Pro SP2 installation. This smaller image gives us better utilization of system resources (i.e., memory). Simple comparisons using the Task Manager show that the number of processes running and the memory required, although varying, is half that of the standard XP Pro installation.

Support Tools

While looking at OS choices we also looked at tools that would allow us to manage our PCs and the OS that are run on them. As we had not decided on an OS and there was still a chance of multiple operating systems, we needed a tool that would not be dependant on the operating systems it controlled. The IT Group at SNS uses Altiris products as its way of managing the workstations used by all of our co-workers and because we are all about using the help of anyone around, we have leveraged the IT Group's experience with Altiris Deployment Solution in their environment and tailored it to our diagnostic PCs. Actually, along with the practicality of using a tool we already had, Altiris was chosen because it provided us with many features including management of many different OS types that now also include MAC OS X.

Altiris Deployment Solution [4] is a client/server PC management tool that includes the ability to create disk images of the PCs hard disk and to distribute that image across the network to other PCs configured to Wake on LAN and PXE boot. This can be done to one or more PCs simultaneously. With our current XPE OS, this process, regardless of the number of machines being configured, takes about 15 minutes. The Altiris software also allows us to deploy software, run specific scripts, and configure each device from one console. Altiris gives us the ability to take a PC from its box and, assuming the supplier has set up the bios correctly, set it in a rack where it will boot and receive its appropriate image.

Figure 1: Deployment System makes it easy to go from Box to Rack.

Windows Server is required for the Altiris software. We have installed it on a dual processor 2.8 GHz server with 2 GB RAM and 140 GB of storage for disk images. This configuration exceeds initial hardware requirements and allows for future growth. We have also decided to use SQL Server in support of the Altiris software. MSDE, the MS free version, did not provide us with the manageability that we required. We have installed SQL Server on another HP DL380, which is configured the same as our Windows server. We have configured this system with three mirrored disk sets: data, indexing, and OS.

Network security is important to the operation of both the accelerator and to these devices. The Diagnostics Group cannot allow others uncontrolled access to the network or these machines. The SNS Controls Group has implemented NetReg, a process that requires all devices that use the network to go through a registration and scan procedure. The use of NetReg along with a comprehensive security plan ensures that the network does not have insecure devices attached to it. Both NetReg and our security plan call for scans of devices prior to being attached to the network to make sure they are virus-free and properly patched.

Currently our backup strategy is a simple one. Altiris allows us to capture machine images and, because of the small number of images deployed, we can do this with relative ease as a job or as part of the initial deployment process. The images are stored on a RAID 5 Storage Array and backed up to tape. We also use the SNS Oracle database to capture specific configuration settings and files. So Oracle, being backed up, provides us with a historic backup of PC settings. We also make use of the database during the boot process. A Java application, using the JDBC thin client, queries the database to retrieve the PC's appropriate configuration settings. This allows us to minimize differences in our PCs and to keep our software somewhat generic.

BENEFITS OF PC USE

Of all the benefits that our PC deployment gives us, two stand out as most important. First, the use of the Microsoft Windows gave us an operating system that was well matched to our developers and the hardware that they required. Second, the implementation of a field replaceable IOC or NAD removed the need for a complete suite of deployed devices having to be shut down for upgrades or hardware failure on one IOC. Both benefits were enhanced by others and proved to be crucial to the success of our deployment.

Windows in General

One advantage of deploying with a Windows based operating system was that it was exactly that. Our partners around the world have delivered to us their various applications and supporting drivers etc. ready for use in deployed systems. No conversion of code was necessary. The software that ran the bench test equipment is the same that runs the beam line devices. This allowed ORNL staff to concentrate on other issues. This was also advantageous for partners as they visited the ORNL site. Partners who looked to modify their code or to see modifications made by others found the code to be exactly as they knew it. Diagnostics developers have created modifications and known exactly how the code should and would be stored on the PC.

NAD

As commissioning continued, we constantly tried to get the most out of our devices. To do this, we have had to make changes to applications running on the PCs that were in use by commissioners. The NAD concept allowed developers to take one instance of a device and implement a new software design. This work was made known to the Operations Group and allowed for in-situ testing before mass deployment. Upon a successful modification, we were able to distribute the changes without disturbing operations.

Also, any hardware failure in the PC or in supporting electronics caused only a single instance of the beam line device to be lost. And if we did lose a PC, down time for that one instance was minimal because of our ability to swap in a new PC and apply a new image. And while making considerable improvements to deployed systems, the NAD implementation has ensured that beam down time attributed to the diagnostic group stays at zero.

XPE

As we started learning XPE, we probably thought about quitting at least once an hour. This desire waned as we learned how to use the developer tools and began to really take advantage of the work that had already been done by others. Creating XPE components for software that is not part of XP Pro can be quite simple or it can be an intern's project, but our unlimited footprint size gives us the ability to eliminate the process of creating complex components (like trying to componentize LabVIEW). So we have a choice and can take the easiest path. Also, XPE has a support community specific to it that can provide answers similar to any Windows problem. There are more than a few Windows users, and we take advantage of that.

Other

MS Windows XP comes with Remote Desktop and this built-in tool has allowed us to control all of our IOCs from our control room or development areas. Also, the remote desktop functionality can be used with the "rdesktop" command available on our Linux work stations.

Other issues that are not generally appreciated include the fact that our PCs do not weigh as much as a VME crate, and this allows ease of PC installation.

PROBLEMS INCOUNTERED

Configuration Control

The most anticipated problem, configuration control, is common to system engineers no matter what the OS. Controlling the different versions continues to be the biggest issue encountered so far. The use of XPE has given every user complete comfort in the environment and therefore a built-in comfort level when making changes. If it is something he can do on a desktop, he believes he should be able to do it on a deployed system. Of course our system users are our developers so most of the upgrading or modifying is for good reason. But XPE allows for tinkering and means that new versions can be done very rapidly. Our ability to distribute the new software enhances the idea that new versions can be tried whenever desired. So far we have been fortunate in that we have forced engineers to keep their different beam line systems unique to the system. With this fortune, we have been able to keep the number of different deployed configurable systems down to under 20. But coming hardware changes and requirements forced by systems deployed for the SNS Ring will only force a growth in the number of configurations in use.

Also, the benefit of having a Windows operating system tempts developers to work with the latest available software. This increases the risk of incompatibility with other applications, and sometimes forces application upgrades to accommodate new functionality. All of this leads to an increase in the number of valid deployed versions.

Now, with the progress we have made, certain issues arise that seem to be unique to individual instances of the IOCs. This type of error was expected in that we know that the global or system-wide issues would be easy to identify. The unique issues require more time and investigation to rectify. With global problems, one solution corrects many IOCs but now we are dealing with one solution per problem.

Too Many IOCs

The Diagnostics Group at SNS is currently responsible for about 190 functioning IOCs. Ten of these are VME crates that support our Beam Loss Monitors and Neutron Detectors. 180 IOCs are PC-based. Some of the basic problems that we have encountered are some that we expected. We have had to make sure that we had the rack space to house them. Not only are we using more rack space but we are also generating more heat within the racks. We also had to make sure that the support systems the PCs need are in place. Plans and hardware had to accommodate these greater numbers. For example, these PCs needed more outlets, timing ports, and network ports than would be expected from a typical VME setup.

FUTURE PLANS AND EXPERIENCE SUMMARY

Future Plans

Depending on stability of our hardware, we are looking at the idea of a diskless solution. This would allow the NADs to boot up and operate from a unique disk image file located on a networked server or servers. The OS would still be Windows XPE and would allow for quick (about the same as hard disk) boot times. This solution will be considered because of the gain in the ability to deploy new images to the same hardware. Before completely converting to a diskless deployment, diskless operation would provide a reliable backup system in case of hard disk failure. Currently we are evaluating a product called BXP from Venturcom [5] that can give us this ability.

The interaction with the relational database gives us the opportunity to expand the database role and use it to capture and house device images as a whole. Although Altiris provides a simple way to handle images and 300-400 images is not that great of a number over all, the relational database naturally lends itself to be the repository of all data related to a device.

Initial Experience

As SNS recently completed its fifth commissioning run, we have really begun to take full advantage of the tools we have implemented. Because of our approach, our initial use of these tools and their implementation required little learning time. This was particularly true when it came to working with MS Windows and Altiris. Windows is our development platform and previous SNS use of Altiris allowed us to quickly deploy it. We did have to go through a learning curve during the start of our implementation with XPE. Initially, it required some patience, and it took us time to ramp up and use it effectively. Now, however, we have a standard image acceptable to users and developers. Inaccessibility to our PCs could have been a major problem but because of the toolset we have included on our image along with other networking tools, we have easily continued to manage our PCs. The PC deployment strategy emphasizes security, and the PCs have not been subject to any security issues since being deployed.

At SNS support for Windows-based PCs was limited so we had to engage in a public relations battle. The success of our deployed systems has begun to garnish support, but we realize that we have only installed half of the IOCs that are planned and we have only been running for a relatively short period of time. Overall, the decision to use MS Windows and specifically XPE has given us the ability to safely and easily deploy, configure, and maintain SNS diagnostic systems. It has allowed us to use standard administrative tools and allowed our developers to use their development tools of choice. We have been able to use off-the-shelf hardware without concerns of compatibility issues with the OS. The time required to deploy an image is far shorter than using standard installation methods, and everyone appreciates under a time constraint. With much pride, we can say that for a new accelerator, our use of XPE has directly contributed to an unforeseen stability in an unprecedented variety of accelerator diagnostic devices and to the over all success of the SNS Diagnostic Group.

REFERENCES

[1] Work supported by the U.S. Department of Energy under contract DE-AC05-00OR22725.

[2] "Network Attached Devices at SNS" W. Blokland and T. Shea, ORNL, Oak Ridge, TN, USA, M. Stettler, LANL, Los Alamos, NM, USA. DIPAC 2003.

[3] SNS Web site: http://www.sns.gov/.

[4] http://www.altiris.com/products/deploymentsol/.

[5]http://www.vci.com.