

SNS ICS Network Final Design Review Report

September 17, 2001

Reviewers:

- Bill McDowell ANL
- Scott Pinkerton ANL
- Ken Sidorowicz ANL
- Karen White TJNAF
- Mike Turpin ORNL

Presenters & Attendees at Oak Ridge:

- Bill DeVan
- Chuck Fisher
- Mike Turpin
- Susan Hicks
- Ernest Williams
- Dave Gurd
- Coles Sibley
- Saeed Assadi

Agenda:

- OVERVIEW - Bill DeVan
- NETWORK MANAGEMENT AND SECURITY FEATURES - Chuck Fisher
- SOFTWARE FEATURES - Ernest Williams
- RELIABILITY FEATURES - Bill DeVan
- COST & SCHEDULE - Bill DeVan
- OLD ISSUES, NEW ISSUES - Bill DeVan

The reviewers all attended via video conferencing links.

1 Review Committee Charter

Verify that the design proposed will meet provide the necessary networking infrastructure to operate the SNS facility in a reliable manner. Make suggestions where appropriate to improve performance, value, maintainability, reliability, and "upgradeability". Give an opinion on whether or not the ICS network is ready to proceed with construction.

2 Executive Summary

The network infrastructure as proposed will provide the necessary networking infrastructure to operate the SNS facility in a reliable manner. Reviewers suggested that the ICS staff investigate the robustness of using LDAP authentication because of single points of failure problems, the use of a terminal server/rebooter device with a hard coded password and the separation of the diagnostics network and the controls network

The committee feels that the SNS should proceed with the purchase and installation of the SNS controls network. The systems being purchased and installed are flexible enough so that the exact topology of the network can be easily changed if testing determines that there is a problem.

3 Detailed comments on the Presentations

3.1 OVERVIEW –PRESENTED BY BILL DEVAN

The reviewers questioned the term “out-of-band management” when all communication links run over the same switched networks. It was noted that some of the redundancy features such as two separate routs for communication cables have been dropped due to cost considerations. Questions also arose concerning the use of shielded CAT 5 cables to equipment in areas of high electrical noise. The committee recommended that standard CAT 5 be used with a fall back position to fiber. There was no concern about having the PLC programming effort use the ICS network.

There is probably some RF info needed by Cryogenic controls. [Therefore if cryo network connection to rest of world is severed, problems with the cryo system may arise]. Redundant network connections between cryogenic controls and the RF controls might be needed.

3.2 NETWORK MANAGEMENT AND SECURITY FEATURES - PRESENTED BY CHUCK FISHER

There was quite a bit of concern about the use of the IOC reboot devices that use a hard-coded IP address and password. SNS should plan on intrusion detection devices to mitigate this risk.

3.3 SOFTWARE FEATURES - PRESENTED BY ERNEST WILLIAMS

Experience with LDAP has shown that it is a single point of failure. IOC names, host names, other computer accounts for operations should be programmed in the “old- fashioned” way. Otherwise loss of LDAP can kill everything.

The presentation showed that back-up service happens “out of band”, but this isn’t really true as it is just another VLAN on the same set of switches.

The presentations failed to show a radius server for account management, authentication, etc. Make sure the interface with SNS IT is defined.

A SANs solution for the backup strategy should be investigated.

3.4 RELIABILITY FEATURES - PRESENTED BY BILL DEVAN

The fact that all network equipment is powered by UPS systems was a good practice. Communication rooms are well air-conditioned but the SNS should monitor the network switch temperatures.

3.5 COST & SCHEDULE – PRESENTED BY BILL DEVAN

There was some concern as to whether installation cost estimate was adequate. Several reviewers felt that the installation costs have been underestimated.

4 Individual Reviewers Comments

4.1 MICHAEL TURPIN ORNL

1. The reboot and terminal server device has a hard-coded password -- Since this password will be [very] difficult to change, it should be well protected. This network should be IDS protected.
2. DNS and LDAP servers are included in the network design, but no RADIUS server is included (needed for dial in and VPN). This oversight should be corrected.
3. Temperature monitoring of switches -- Good planning has been done for redundancy (switch fail over, power backup, etc.), but no HVAC redundancy/contingency is mentioned. In lieu of redundancy, temperature monitoring should be included AT A MINIMUM. Most Cisco switches provide this capability via SNMP. HP Openview (or equivalent) and a management workstation should be included for network monitoring.

4.2 SCOTT PINKERTON ANL

1. As a general note - I (personally) have a feeling of surprise over the lack of "A Specs" or other high-level requirement type documents. [Maybe this documentation exists, and I have just

not run across it. Though if it existed - I think that we should have been reading, reviewing, and commenting on all of the those documents - not just the power point slides.]

2. From the preliminary design review - we had made some suggestions about "administrative services" and the need to establish appropriate flow-down requirements early on. In Bill DeVan's set of slides "Old Issues/New Issues" pg. 2 - he indicates (on the fourth bullet) that they have successfully addressed administrative services E.g. DHCP. I question what the high level requirement really is.

3. With that in mind it seems difficult to imagine how we ended up with the "Rabbit 2000" single-board computers that require their passwords and IP addresses to be hard-coded. This seems like a bad decision here - the long-term "life cycle" costs of living with these limitations might exceed the costs of re-designing this item now. If they are "semi-custom" (some amount of internal development), then it seems like we should put a bit more effort in here. [Contrasting this to devices like PLCs, were it does not seem practical to try and re-design for certain limitations.]

4. Staying with the concept of the "administrative services" (which we commented on from the preliminary design review) - I would be interested to see how the high-level requirements read on the topic of "name services". Will flat file based host tables be the primary method of management, will traditional DNS (Bind type implementations) be the primary method of management, or will LDAP be the primary. I agree with Steve Lewis's comments that today LDAP might represent too much risk for this to be the primary "name service". Overall it is not clear to me - what will be the primary method of implementation and what will be backup/secondary.

5. Still on "administrative services" - Mike Turpin pointed out the need to identify a radius server to support authentication. For the sake of completeness (and the benefit to accurate ETC costs) I would suggest identifying a requirement for an IDS director/management console. Same with a network management console (E.g. HP OV, Cisco Works 2000, something). Also what about a centralized syslog server - again for many reasons (security being high on the list) it is becoming more important to retain the log files for considerably longer than we had in the past. [This concept should not be lost wrt to DHCP services. Where do we document how long the records for the DHCP IP assignments should be retained? If you had a security incident and wanted to know who/what was using an IP address last week/month/year - will we always be able to access that data?]

6. I thought the idea of defining enterprise back-up system and services up front was a very good idea. [Excellent info to go into an A spec.] Though there was an odd reference to "out-of-band" networks. This was discussed during the review - how/where does the documentation get fixed?

7. Regarding "reliability and maintenance features" - I still don't feel comfortable on what the A spec requirements are (or should be here). As such, how much downtime is acceptable (2.3 hours/10,000 or 1.4 hours/10,000). Are there any requirements (or design goals) that we should be striving for vs. just lower cost? As such, other than the cost argument - it is difficult to advise on using a single supervisor in the Catalyst 6509's; it is difficult to comment on the use of a single large fiber optic cable to the comm. rooms vs. two cables routed separately. It would be in a different WBS section - but I was unable to read about power distribution issues (assumptions for the comm. rooms) and how that should "flow-down" into the control system design. Too me, again these are things that would show up in an A spec that defines your working interface (or possibly in an Interface Control Document (ICD)).

8. Regarding "cost and schedule" - the current ETC numbers just look too low (in my humble opinion). I was concerned on more than just the cable installation numbers. I don't see enough effort for things like a) writing A specs (requirements documents) b) working on ICDs c) overall integration support of the control system with other portions of SNS.

9. Some effort obviously had been done in developing a security plan. Unfortunately, it still seems too much "in development" for this being a final design review. I would strongly encourage that additional effort be spent on the security plan to drive out the "derived requirements" for things like an IDS system. Or for something like - non re-usable passwords must be used for certain devices like (core routers, switches, authentication server ala radius and LDAP, and DNS servers, etc.).

4.3 KAREN WHITE TJNAF

1. There must be a highly reliable way for information about the RF load to get to the Cryogenics system. The information as presented showed that the Cryogenics system can operate in the absence of all other accelerator systems. The RF information is most likely critical to the proper functioning of the Cryogenic Systems and there should be some mechanism in place to ensure this can be transmitted in the absence of the full normal network operations.
2. One viewgraph said that PLC programming would take place on the Controls Network. There was some discussion or comment concerning this. I don't really view this as a problem since it is unlikely the PLC programs will normally get changed during machine operations.
3. There was some discussion of LDAP as a single point of failure and a suggestion that certain machine critical accounts should be maintained separately. I think if one analyzed this, the number of accounts could be in the hundreds and the information would need to be maintained on (at least) dozens of computers. I am not very familiar LDAP, but with DNS, we run alternate servers on other machines that take over in the event of a failure of the primary DNS server. Something like this would seem a better solution than maintaining separate lists of many accounts on many computers.
4. Backup Strategy – The material noted additional network cards in each workstation for use by the backup system. Is it actually necessary to backup each workstation on the network? I am guessing there will be well over 100 workstations. Planning to backup each of these machines implies there is unique information contained within. One might consider a central file server for each segment (e.g. Cryogenics, Accelerator, Target, etc.). Then, since each workstation gets files from the server, it is easier to ensure everyone has the "right" copy of a file, and only the file servers need to be backed up.
5. The account management strategy needs some work, but this might more properly fall under "Security" than "Network". It was noted that management of Rabbitt with hard-coded passwords will be difficult, especially when people leave. One could argue that people first need to access another machine (where accounts can easily be disabled) before accessing the Rabbitt. Connections to this device could also be logged so it is easy to tell who is using it.
6. The separation of the Accelerator and Diagnostics networks probably has implications for some of the more complex applications that will eventually be developed. For example, I can imagine a Fast Feedback System that would need to access both of these networks and would need to operator at very high speeds meaning the planned accommodations for passing data between the networks May not be adequate. This is probably only applies to a small number of applications, but could pose a real problem.
7. In the last section, we noted that the labor allocated for technicians and electricians may be on the low side. Regarding the actual charge to the committee, I do not see any issues that should hold up network hardware procurements. The issues that need to be worked out seem to be policy or system management issues and do not imply the hardware configuration needs to change.

4.4 KEN SIDOROWICZ ANL

The presentation on the network Sniffers mentioned moving cables to connect the sniffer to subnets. I recommend that a sniffer switch be ordered to avoid this problem. I also note that the design has many small network switches. This will cause the staff extra work when it comes time to upgrade the firmware on these switches. (usually two to three times a year) Every switch will need to be physically visited to perform these upgrades.

4.5 **BILL McDOWELL ANL**

Architectural details of the interface between the Diagnostics and Controls are not clear. It appears that the Diagnostics group has elected to use computer hardware that is significantly different than the rest of the project is using. This difference in both base hardware (PCI vs. VME-VXI) and cpu type (Intel vs. power pc) while possibly saving construction money will lead to increasing the long-term support costs. These systems will tend to have different problems and require different knowledge bases and possibly different staff members for solution. It is definitely too late to change his direction, so we have to make the best of it.

The network plans show a separate subnet for Diagnostics (using vlans)but a detailed reason for separating the networks was not give. It should be noted that commissioning teams will need access to all data in the form of EPICS process variables. The commissioning team should be able to create scripts and application programs without working through the Diagnostics group staff. To do this they will need access to all the raw and processed data.

5 **Conclusions & Recommendations**

The network infrastructure as proposed will provide the necessary networking infrastructure to operate the SNS facility in a reliable manner. The committee recommends that the SNS proceed with the purchase and installation of the SNS ICS network. The systems being purchased and installed are flexible enough so that the exact topology of the network can be easily modified if testing determines that there is a problem exchanging data between Controls and Diagnostics. The use of LDAP should also be further investigated because of a possible single point of failure problem.