

**Brookhaven National Laboratory
The STAR Experiment's Data Acquisition System
(BNL-SDAS)
October 2006**

- I Security Plan*
- II Contingency Plan*
- III Risk Assessment*

BROOKHAVEN NATIONAL LABORATORY THE STAR EXPERIMENT'S DATA ACQUISITION SYSTEM (BNL-SDAS) OCTOBER 2006..... 1

I BNL STAR EXPERIMENT'S DATA ACQUISITION SYSTEM SECURITY PLAN4

1 IDENTIFICATION4

1.1 NAME/TITLE 4

1.2 TYPE 4

1.3 SECURITY CATEGORIZATION 4

1.4 OPERATIONAL STATUS 4

1.5 RESPONSIBLE ORGANIZATION 4

1.6 CONTACTS 5

1.7 RELATED LAWS/REGULATIONS/POLICIES 5

2 DESCRIPTION.....6

2.1 ENVIRONMENT..... 7

 2.1.1 *Physical Environment*..... 7

 2.1.2 *Physical Access* 8

3 MINIMUM SECURITY CONTROLS9

3.1 ACCESS CONTROLS (AC-1) AND ACCOUNT MANAGEMENT (AC-2) 9

 3.1.1 *Remote Access (AC-17)* 10

3.2 AWARENESS AND TRAINING (AT) (OPERATIONAL CONTROL) 10

3.3 AUDIT AND ACCOUNTABILITY (AU) (TECHNICAL CONTROL)..... 10

3.4 CERTIFICATION, ACCREDITATION, AND SECURITY ASSESSMENTS (CA) (MANAGEMENT CONTROL)..... 10

3.5 CONFIGURATION MANAGEMENT (CM) (OPERATIONAL CONTROL) 10

3.6 CONTINGENCY PLANNING (CP) (OPERATIONAL CONTROL)..... 11

3.7 IDENTIFICATION AND AUTHENTICATION (IA) (TECHNICAL CONTROL)..... 11

3.8 INCIDENT RESPONSE (IR) (OPERATIONAL CONTROL) 11

3.9 MAINTENANCE (MA) (OPERATIONAL CONTROL) 11

3.10 MEDIA PROTECTION (MP) (OPERATIONAL CONTROL)..... 11

3.11 PHYSICAL & ENVIRONMENTAL PROTECTION (PE) (OPERATIONAL CONTROL) 12

3.12 SECURITY PLANNING (PL) (MANAGEMENT CONTROL)..... 12

3.13 PERSONNEL SECURITY (PS) (OPERATIONAL CONTROL)..... 12

3.14 RISK ASSESSMENT (RA) (MANAGEMENT CONTROL)..... 12

3.15 SYSTEM AND SERVICES ACQUISITION (SA) (MANAGEMENT CONTROL)..... 12

3.16 SYSTEM AND COMMUNICATIONS PROTECTION (SC) (TECHNICAL CONTROL)..... 12

3.17 SYSTEM AND INFORMATION INTEGRITY (SI) (OPERATIONAL CONTROL) 12

II BNL-SDAS CONTINGENCY PLANS13

III BNL-SDAS RISK ASSESSMENT14

1 IDENTIFIED RISKS.....14

2 THREAT IDENTIFICATION14

3 RISK ANALYSIS14

4 CONTROL RECOMMENDATIONS.....16

Responsibility for this information subsystem and its operation as described in this document is accepted by:

_____ **Date:** _____
System/Application Owner
Jérôme Lauret

_____ **Date:** _____
Computer Protection Program Manager
Keith Lally

_____ **Date:** _____
Information Technology Division Director
Thomas Schlagel

_____ **Date:** _____
Research Enclave Owner
Peter Bond

I BNL STAR Experiment's Data Acquisition System Security Plan

1 Identification

1.1 Name/Title

This document describes the Brookhaven National Laboratory, Relativistic Heavy Ion Collider (RHIC) Solenoid Tracker At Rhic (STAR) experiment's Data Acquisition system. It is referred to as the STAR Data Acquisition System or BNL-SDAS.

1.2 Type

The BNL-SDAS is an information subsystem in the BNL Research Enclave.

1.3 Security Categorization

The Security Categorization of BNL-SDAS as defined by NIST FIPS 199 is
SC BNL-SDAS = {(confidentiality, LOW), (integrity, LOW), (availability, LOW)}.

The information on BNL-SDAS is experimental raw data which will result in scientific research that will be published in the open literature. No harm will result from loss of confidentiality. Compromised data integrity would require in depth knowledge of the data format which would be detected by the data reader embedded checksum. Metadata information is backed-up and/or replicated; the data only transits through this system. Integrity loss would lead to physics results that are inconsistent with other data sets and has a high chance to be caught by either our online or offline Quality Assurance procedures. Unavailability of the system may delay data acquisition but is within acceptable time loss generally accepted by collaboration of this nature and integral part of normal operation expectations. Unavailability is mitigated by components model in the system: they are redundant and/or failover.

1.4 Operational Status

BNL-SDAS is in the operational phase of its life-cycle.

1.5 Responsible Organization

Brookhaven National Laboratory
P.O. Box 5000
Upton, NY 11973

1.6 Contacts

Title	Name	e-mail	Telephone
Information System PI	Timothy Hallman	hallman@bnl.gov	(631) 344-7420
Computing & Security Lead	Jérôme Lauret	jlauret@bnl.gov	(631) 344-2450
System Engineer	Jeffery Landgraf	jml@bnl.gov	(631) 344-7967
System Engineer	Jon Engelage	JMEngelage@lbl.gov	(510) 486-4827 (631) 344-5371
System Engineer	William Waggoner	BillWaggoner@creighton.edu	(402) 280-1791 (631) 344-7803
System Engineer	Michael DePhillips	dephilli@bnl.gov	(631) 344-2499
System Administrator	Wayne Betts	wbetts@bnl.gov	(631) 344-5795

1.7 Related Laws/Regulations/Policies

BNL-SDAS is subject to the laws and requirements described in the Brookhaven National Laboratory Cyber Security Program Plan (CSPP).

2 Description

The BNL-SDAS is an online system dedicated to support the STAR apparatus data collection, monitoring and quality assurance during RHIC runs. This non-classified data is essentially transient to the system and shipped to either the RHIC Computing Facility (RCF) Mass Storage System (HPSS for the raw data) or to offline databases (for some MetaData) via controlled data migration or database replication.

The BNL-SDAS system is composed of 3 main components. All machines are behind a firewall maintained by ITD and/or inside an un-routable (private) network within that firewall. The whole installation is confined to the STAR complex (Building 1006). The components are:

- Slow Controls – this component is comprised of Front-End VME-based processors that run the VxWorks Operating System and a small number of Linux PCs. The PCs provide the OS and configuration at start-up for the VME devices, poll the VME processors for data, and provide a user interface via EPICS components. The VME processors typically read, configure or control various parameters of detector hardware components. They do not typically collect physics data per se, but they are the source for some MetaData and calibration parameters that may be stored on other systems. Those VME devices are located within the firewall or the un-routable network.
- Core STAR trigger/DAQ - this component consists of readout electronics (thousands of boards) detector mounted or housed in VME or CAMAC crates controlled by diskless CPU's and connected together using Ethernet, serial lines, home-grown node to node fiber protocols and a Myrinet network. Computing nodes on the DAQ network consist of 50 diskless MVME processors running the VxWorks operating system, used to read data from the electronics and 10 Linux machines used for control, monitoring, quality assurance, event building, event buffering and data transfers to the RCF data storage facility. Data arrives continuously during RHIC Physics Beam operations at rates as high as 100MB/s. The event building machines provide a combined data buffer of 4 TB, representing roughly 10 hours of Au-Au beam operations with central trigger. The software consists of hundreds of thousands of lines of highly specialized code. VME crates within this category are all on the un-routable private network.
- The database system – this component is primarily a (Meta) data taking system. These data include condition and monitoring data from the detectors and the collider as well as event, file and run tag information from the data acquisition system. The exception to this is an aggregate series of information that compiles vital run information and displays it in real time on the web. The detector and collider data are streamed in to the database tables using a suite of C++ based daemons that poll STAR's slow controls system (EPICS). Each daemon is specific to one data-source (e.g., a specific detector) and writes to a specific database. Due to the specificity of the daemons and since they are compiled C++ code they are fast, have a small footprint and are independent of each other, thus creating no dependencies. The Run, File and Event tags are written directly into the databases, using the MySQL C API, from the data acquisition. Migration of the data is done via either native database replication mechanism or a set of processes build from

home grown code aimed to verify the data sanity, aggregate and/or average values, compare to previous time snapshot and inserting a new entry to a normalized table.

There are about 20 additional interactive desktop machines for the shift crews and support personnel to use in the Control Room, which display detector status and controls, Run Control and QA monitoring information. The network is 130.199.60.0/23. The buffer boxes and a few designated machines are dual-homed on the non-routable "HPSS network" 192.168.10.0/24, which provides a dedicated high-bandwidth connection to the RHIC Computing Facility for long term data storage.

2.1 Environment

2.1.1 Physical Environment

The BNL-SDAS computers are located in the principal building of the 1006 complex at the 6 o'clock position of the RHIC ring. BNL-SDAS components reside in several distinct areas, known as the Interaction Region (IR) (also known as the Wide Angle Hall) where the actual detector systems are located during data taking, the Assembly Building where the bulk of the detector may be moved during maintenance periods, the Gas Mixing Room, the DAQ Room and the control room. The DAQ Room and Control Room contain the majority of the computers and online data processing and transferring components. STAR shift crews run the experiment mainly from the control room. Two small modular buildings (1006C and 1006D, often referred to as the "trailers") are used by visiting researchers and support personnel, but do not include components of the BNL-SDAS and any computer systems in the trailers fall within the general BNL Research Enclave. Trailers are lockable, but during the run are unlocked, and manned 24x7 by members in good standing of the collaboration. Unrecognized visitors to the trailers are challenged by SDAS personnel as a matter of procedure.

Key components requiring data integrity (database essentially) are powered through several UPS systems with the capability to operate on battery power for at least 20 minutes without line power. This allows sufficient power surge protection and more importantly, enough time to safely shut down the servers as needed in the event of power failure.

The fire suppression and interlock system is part of the general safety system for RHIC experiments. The STAR Global Interlock System (SGIS) is a primary physical safety system for hazards to both personnel and detector systems, and as such undergoes intensive annual reviews and certifications. It is configured to power down significant portions of STAR, including computing equipment located in the DAQ Room and/or IR, upon detection of water, smoke, fire, flammable gas, or other hazards.

2.1.2 Physical Access

During a RHIC Run, the Ring Road that provides vehicle access to the 1006 complex is posted "Radiation Controlled Area" starting at the 5 o'clock position. Only authorized and properly trained personnel are permitted to access the 6 o'clock position. During a Run, the control room is occupied around the clock, (i.e., 24 hours a day, seven days a week including holidays) with the shift and other support personnel. Shift crew composition is verified according to the procedure described in section 3.1. Outside of a Run, all doors remain locked.

3 Minimum Security Controls

The sections below provide information for the implementation of controls from each of the 17 families of low risk security controls from NIST SP 800-53 that applies to BNL-SDAS where they differ from the general Research Enclave description.

3.1 Access Controls (AC-1) and Account Management (AC-2)

There are a number of operational accounts which are used to run and monitor the experiment, such as one to run the actual data acquisition, one to remotely control the High- and Low-Voltage systems of the detector system components, and one for the Forward TPC sub-system. Those accounts typically remain logged in on the control room consoles for the duration of a Run. The access passwords are known by the personnel listed as contacts and the ongoing Shift Leader. The shift personnel inherit the session from the outgoing shift person in the respective position and do not need to log in.

In some cases, an expert may require access to a special account on certain machines in order to fix a problem or monitor the behavior of a component or a detector system. We grant and control access by way of two-factor authentication (ssh keys), which the expert in question generates and maintains on an RCF computer. The access therefore follows the standard BNL account approval while providing a manageable way to grant, control, and document access of individuals.

For databases, user privileges are limited to general reader groups for selection of data. Inserts of records is reserved for and limited to approve processes via Writers accounts of limited access. Administrative tasks can only be conducted on the console of the local hosts which are password protected.

The shift personnel are Members in Good Standing in STAR and hold computer accounts at the RCF. At any point in time, the Shift crew is composed of 4 STAR collaborators, as well as from 0 to 2 trainees. There are three Shifts per day, and a Shift assignment is 8 days. There is one day of overlap between the outgoing and incoming Shifts to facilitate updated Training for the current run year. In details, the Shift Crew is composed of a Shift Leader, two Detector Operators, and a general Shift crew position. The Shift Leader and Detector Operators must have previously obtained on the job training at STAR, which results in written evaluations of their performance submitted to a QA Board, and the determination by the QA board that the individuals be certified to serve in these positions. Lists are kept of all certified Shift Leaders and Detector Operators by the shift coordinator. The personnel on Shift for any given day and Shift during the run are posted on a Collaboration Web site and all receives basic STAR training including proper conduct with respect to the STAR computing infrastructure.

Each Shift Leader is expected to “hand-shake” with his crew when first in place and each shift assignment period requires a sign-out and sign-in form allowing to pass on from two consecutive crew a list of ongoing issues and/or configuration information for the run period. The “shift change-over” form is electronic and is kept in an Electronic shift log as well as sends a summary to a mailing list.

3.1.1 Remote Access (AC-17)

Personnel listed in the contacts and experts for various components of the STAR Detector can remotely access to the BNL-SDAS facilities. In order to use the already existing access controls and account management infrastructure that is in place for the RHIC Computing Facility, we use the exiting RCF SSH gateway machines to access the BNL-SDAS computers. This is the only access method from off-site or the external wireless network. There are no firewall conduits from the external Internet to BNL-SDAS.

3.2 Awareness and Training (AT) (Operational Control)

BNL-SDAS conforms to the requirements of the CSPP. Remote access to experts is granted via RCF account and `ssh` gateways. This requires all users to complete BNL Cyber Security training (see <http://training.bnl.gov/>, course GE-cybersec) before receiving a RCF account.

3.3 Audit and Accountability (AU) (Technical Control)

BNL-SDAS conforms to the requirements of the CSPP.

3.4 Certification, Accreditation, and Security Assessments (CA) (Management Control)

BNL-SDAS conforms to the requirements of the CSPP.

3.5 Configuration Management (CM) (Operational Control)

Systems on the 130.199.60 subnet conform to the requirements of the CSPP as members of the Research Enclave. Accordingly, Linux and Windows systems are regularly patched with vendor-supplied patches. There are several Redhat Enterprise Linux systems that use the BNL Redhat Enterprise Network Satellite Server, while systems with the predominant Linux distribution (Scientific Linux) typically receive updates via nightly yum cron jobs that access the Physics Department's SL mirror site. ORDO is being deployed on all nodes capable of running it. Windows systems have TrendMicro OfficeScan and SUS installed (with migration to SMS).

There are certain networked devices that are not configurable in the conventional sense of a typical networked computer, because their OS is not intended to be interactive other than the most basic functionality or, in some cases, the OS is not upgradeable for technical reasons or is no longer vendor-supported. These devices include networked power supplies, serial console servers and other hardware interfaces. They may rely on protocols (e.g. telnet) that should not be widely exposed. These devices are placed on a dedicated un-routable network with a limited number of dual-homed gateway systems from which authenticated users are able to access them. The shift supervisor would immediately note any changes to the power status. Additionally, there are hardware safety locks that prevent over-power to devices.

The databases are configured identically on two nodes. Both nodes are built using a series of configuration scripts and contain a complete set of data. If one node goes down the system will switch to the second immediately; if both nodes fail, the database can be brought back-up on a spare node using configuration scripts and data from back-ups within a few hours.

3.6 Contingency Planning (CP) (Operational Control)

The BNL-SDAS Contingency Plan is included in Section II of this document.

3.7 Identification and Authentication (IA) (Technical Control)

The use and access to special (shared) accounts is described in section 3.1 of this document. Console accesses are locked using standard screen saver/locker mechanism. BNL-SDAS otherwise conforms to the requirements of the CSPP.

3.8 Incident Response (IR) (Operational Control)

BNL-SDAS conforms to the requirements of the CSPP.

3.9 Maintenance (MA) (Operational Control)

BNL-SDAS conforms to the requirements of the CSPP.

3.10 Media Protection (MP) (Operational Control)

BNL-SDAS conforms to the requirements of the CSPP.

Core file systems are backed up over the network using ITD's EMC Networker backup system, which provides weekly full backups and daily differential backups. This provides offsite storage as well as rapid "self-service" restoration capability in most circumstances, such that an administrator or subsystem expert can complete the restoration of individual files or even entire file-systems as needed without involvement of ITD personnel. Backup lifetime is one month, though for extended shutdown periods (typically three to five months each year) an "end-of-run state" full backup may be stored longer as deemed appropriate. "Core" file systems are identified by consultation between the BNL-SDAS technical managers and subsystem experts.

A CVS repository is made available for subsystems to store their individual specialized software, from which the software could be deployed on other nodes if the original were unavailable for some reason. The primary meta-data databases are redundantly stored on two database servers. Additionally, subsystem managers use other backup systems of their choosing to fit the specific parameters of their subsystems, such as periodic "burning" to optical media and local disk to disk backups providing additional snapshots.

3.11 Physical & Environmental Protection (PE) (Operational Control)

BNL-SDAS conforms to the requirements of the CSPP.

3.12 Security Planning (PL) (Management Control)

BNL-SDAS conforms to the requirements of the CSPP.

3.13 Personnel Security (PS) (Operational Control)

BNL-SDAS conforms to the requirements of the CSPP.

3.14 Risk Assessment (RA) (Management Control)

BNL-SDAS conforms to the requirements of the CSPP. The BNL-SDAS Risk Assessment is included in Section III below.

3.15 System and Services Acquisition (SA) (Management Control)

BNL-SDAS conforms to the requirements of the CSPP.

3.16 System and Communications Protection (SC) (Technical Control)

BNL-SDAS conforms to the requirements of the CSPP. A local un-routable network is created via a layer of switches and 5 gatekeeper machines allowing communication from the 130.199.60.0/23 subnet. Network devices such as remote power supplies are accessible from the un-routable network which isolates those devices from scanning or direct access. The routable network and its firewall layer is as described in section 2. Inclusive access to the online Web server (port 80/8080) , ssh access is allowed only through two layers of gatekeepers (a login to the RCF gatekeeper is necessary before accessing the 60 subnet).

3.17 System and Information Integrity (SI) (Operational Control)

BNL-SDAS conforms to the requirements of the CSPP.

II BNL-SDAS Contingency Plans

The impact of a complete loss of the BNL-SDAS system as a whole is significant to STAR although its likelihood is low. The only conceivable cause for such a complete loss is a catastrophic event such as a fire, a flood, an explosion or similar major disaster. A complete loss of the BNL-SDAS hardware during a RHIC Run would mean the end of data taking for the Run since the replacement turn around would involve a procurement process which would take weeks to complete its cycle.

The impact of losing an individual machine is low.

Spares for the VME processors, terminal servers, remote power switches, and myrinet are available and can all be changed and reconfigured within 30 minutes. Most key components run on off the shelf commodity hardware for which spare parts are always available. Downtime due to computer problems is relatively rare however. While the BNL-SDAS system as a whole must function, no single computer is critical by itself; its role can be taken over by another machine within minutes by a failover / load sharing mechanism or a quick turn around of hardware replacement typical of such operation. In normal operation, a defective node is replaced and the load balancing is restored without impact on operation.

Of the currently 4 buffer boxes, 1 could fail without performance degradation. A failure of 2 would cause performance losses only at the highest RHIC luminosities.

The database server has a real-time redundant backup using a virtual IP address. Upon failure of one machine the primary machine can be switched via script instantly. One set of the current run data is backed up nightly by reading a snapshot using native MySQL backup procedure onto a third machine. Upon failure of the two machines the database can be rebuilt in under an hour's time frame. The database snapshots are also backed up nightly via the ITD Legato backup system.

Significant planning was done to prepare for the scenario of losing connectivity to the HPSS Mass Storage system, or experiencing a failure of the Mass Storage system itself. Depending on the severity of the loss of connectivity, the decision to take corrective action will be made by either SDAS staff or BNL senior management. The connection to HPSS is a dedicated network of 2 redundant fibers that provide a combined bandwidth of 2000 Mbit/s. A loss of connectivity would immediately trigger an alarm. The primary purpose of the buffer boxes is to provide a buffer to ride out short service interruptions of the HPSS system. The 4 buffer boxes have a combined capacity of 4 TB. The maximum data taking capacity from the STAR DAQ is 100MB/s (a more typical value in a Au-Au run of RHIC is 60 MB/s). At maximum rate, we could therefore continue data taking for as long as 11 hours. With the typical rate of 60 MB/s, this would equate to 19 hours. In addition, RHIC does not deliver beam continuously but in discrete fills that last about 6 hours, with about 1.5-2 hours setup time in between which further expand the time availability for recovering from an HPSS system or the network problem.

III BNL-SDAS Risk Assessment

1 Identified Risks

The threats, vulnerabilities and risks identified in this section are specific to the BNL-SDAS Information Subsystem. There are also site-wide threats, vulnerabilities and risks that are described in the BNL Cyber Security Risk Assessment and Mitigation document.

2 Threat Identification

A threat is the potential for a particular threat-source to successfully exercise a particular vulnerability. A vulnerability is a weakness that can be accidentally triggered or intentionally exploited. A threat-source does not present a risk when there is no vulnerability that can be exercised.

There are no threat sources that have not been identified in the BNL Threat Assessment.

There are no motivations and threat actions that have not been identified in the BNL Threat Assessment.

3 Risk Analysis

1. THREAT: Some devices do not allow or perform patching, scanning or fail to meet other research enclave controls because it is not feasible due to operational need or technical issues.

Likelihood: Moderate

Impact: Moderate – Exploit of un-patched systems could disrupt operations and put BNL-SDAS systems at risk.

Unmitigated Risk Level: Moderate

Mitigation: This category is mainly composed of “dumb” devices such as remote power supplies to the exceptions of a few Linux systems making use of direct IO. This is known to work only in specific kernel versions: in such cases, the node could be rebuilt from scratch and are redundant failover configured. STAR personnel additionally isolate such devices from the network as described in the Minimum Security Controls section above.

Residual Risk: The following residual risks remain

[Low] If an exploit of a gatekeeper was successful, and there was a subsequent intrusion, there would be some disruption which would be immediately visible and require immediate attention from the shift crew, including notification to Cyber Security.

SUMMARY OF RESIDUAL RISK – LOW

2. THREAT: Unlocked consoles with open sessions or applications running on any single unattended control component.

Likelihood: Low

Impact: Moderate – This threat would only disrupt the sub-system component of the BNL-SDAS that was specifically attacked.

Unmitigated Risk Level: Moderate

Mitigation:

The consoles and computers are in an access-controlled area as described in section 2 above. Computers sessions in the control room are further automatically locked on activity (or lack thereof) timeout.

Residual Risk: The following residual risks remain

[Low] Only after compromising the physical access controls and or disregarding all training material could the threat become a risk.

SUMMARY OF RESIDUAL RISK – LOW

3. THREAT: Some of the VME crates from Slow control are visible from within the firewall and therefore exposed to brute force attack from other devices present on within this firewall.

Likelihood: Low

Impact: Moderate – This threat would be immediately detected and affecting a sub-system. Values can be reset immediately.

Unmitigated Risk Level: Moderate

Mitigation:

All VME crates are being monitored. Alarms goes off shall the set value be tempered independently of the Slow Control front end interface. If threshold settings are exceeded, the crates shut down, leading to an immediate detection. QA monitoring in place further detects configuration changes. In addition, are only allowed to access the network within the firewall, devices which hardware address are pre-known, no other device can dynamically acquire a valid IP.

Residual Risk: The following residual risks remain

[Low] Repeated scanning of VME crates without affecting the system were reported by the community as impacting the system's response.

SUMMARY OF RESIDUAL RISK – LOW

4 Control Recommendations

The likelihood of all identified threats to BNL-SDAS is low, while the impacts are no more than moderate and limited to individual components. There is low risk from all identified vulnerabilities and no additional controls are recommended at this time, outside of what is recommended for the Research Enclave as a whole.