

Notes for a LIGO Computer Security Policy Document

As I will describe in the following, there are a number of issues that motivate a new look at computer security for LIGO. Albert Lazzarini asked me to examine these matters and suggest a course of action.

Implementation Plan or Policy Document

A LIGO computer security plan was prepared a couple of years ago in response to an action item in an NSF computing review. This plan in draft form (I believe) satisfied the subsequent NSF review committee because it indicated that LIGO was paying attention to computer security. The plan was prepared after extensive discussions among the key technical players at the LIGO Lab sites. Its main strength, a comprehensive and thoughtful description of the LIGO network and computing environment as it existed and as it was then planned to evolve, is also its main weakness. This description can be interpreted as the plan itself and as such implies that changes beyond the description are not permissible. This means that the rapid responses necessary to adjust with agility to changes in the LIGO scientific environment or the external network environment are caught between either ignoring the plan or being subject to considerable impedance.

Although some work is going on to modify this document to adjust to changes, the difficulty of doing so is a good indication that such an implementation plan is not a good *high level* document for LIGO. What is missing is a high level LIGO Computer Security Policy document that covers the main considerations in a way that is clear enough that it can be interpreted sensibly in any context but not so detailed and specific that it will be out of date before it is approved. Such a policy document should be intended to survive with only rare evolutionary changes for 5 years or more. The policy may call for more detailed and specific implementation plans to be developed within the various organizational responsibilities that make up the LIGO computing community.

The notes that follow could form the basis for such a Policy document. LIGO has been informed that a computer security discussion will be at the NSF review in fall. Now is an excellent time to prepare a Policy, as this could be the focus of LIGO's presentation on the subject, avoiding a counter-productive extended discussion on implementation specifics. I hope that the following can provide a helpful basis for reaching a consensus on what should be in LIGO's Computer Security Policy.

Scope

The policy document covers all computers and software that touch any LIGO data and all LIGO Laboratory, Tier 1, and Tier 2 computers connected directly or indirectly to the external network.

Goals and Requirements

LIGO's single mission is gravitational wave fundamental scientific research and its single overriding goal is to maximize its scientific output. Computer security for LIGO must be

consistent with this goal and must be designed to address no more than the specific LIGO computer security goals:

1. The primary computer security concern is to avoid disruption of operation or corruption of data.
2. Serious embarrassment caused by defacement of LIGO publicly accessible websites or the use of LIGO computers in criminal activities on the external network can be disruptive to LIGO's science because of the public and government scrutiny that may result. Defense against intrusions that could result in such embarrassment is therefore an important secondary goal of LIGO's computer security.
3. LIGO results are published openly to the scientific community and there is no government classified or export control information involved of any kind. During periods when data is analyzed, the analysis and data are kept from public view in order to avoid widespread publication of incorrect results. This confidentiality by itself requires only a light level of access control, and is a lower priority security goal.

A Chain of Responsibility (and Authority)

LIGO's computing falls under a diverse set of organizational units both within the LIGO Laboratory and at the external institutions that support the Tier 2 Centers. The support for computing within various computing environments such as general desktop computing, data acquisition, and different flavors of analysis computing crosses many of these organizational lines. The responsibility for computer security, and the authority to set security restrictions, at LIGO is a hierarchy of two tiers.

At the highest level, there are two roles that need to be defined and filled. One is an individual at an "executive" level who is a direct report to the Director and Deputy Director. This person establishes the LIGO wide Policy (through the Policy document). He/she has authority over all LIGO computing in regard to computer security issues. Since this role is similar to the Senior Safety Officer, we might name it the Senior Computer Security Officer (SCSO). The SCSO is supported by an individual who oversees the day-to-day implementation, advises and assists others in LIGO on technical matters, and leads the response to intrusions and other serious incidents. This role might be referred to as the Computer Security Manager.

Below the LIGO-wide level, the responsibility and authority for security should fall on those who have to support communities of users or operations as system managers of one or more computers normally located at a single site. This puts the responsibility right where the inherent conflict implicit in computer security lies, where disruptions caused by intrusions or impediments to science caused by excessive computer security restrictions are most acutely felt.

Examples of such "autonomous" computing units are: system manager supported desktop computing at LLO; a single self supported computer on the wireless network at LHO; the LDAS computer system at LLO; the UWM cluster; the PSU cluster; the CDS systems at LHO; the LDAS system and tape library at CalTech; the DMT computer at LHO; a future Matlab environment at LLO; the Tier 2 computer at MIT; the GEO system; etc.

Each autonomous computing unit would have a named Computer Security Officer, normally the person who is also responsible for on site system management. This person would work with users of the unit to maintain a brief written implementation plan, including specific network firewall restrictions at the perimeter of the unit.

This implies a somewhat new concept in network topology for LIGO sites, which I am told is straightforward to implement and is already in the thinking of the LIGO site system managers who have been discussing an updated security plan. Rather than protecting each physical site at the boundary, the Cisco routers would be configured to control traffic in and out of subnets at each site.

Each autonomous computing unit would have its own subnet and its CSO would define the rules for traffic in and out of the unit, whether from off site or from another on site unit. The importance of this is that intrusions within one unit could not spread disruptions to another more restricted unit unless the second unit had permitted that kind of traffic. Thus each unit could balance its protection level where its user community tolerance for disruption corresponds to its tolerance for restrictions. The community will quickly complain if the balance gets out of line.

More about this below

Critical Systems

One class of computing must get special attention. These are systems where an intrusion could cause an unrecoverable and significant loss of science opportunity through disruption or damage to equipment or irretrievable loss of data. The most obvious example of such “Critical Systems”, and probably the only ones at LIGO, are the CDS interferometer control and the corresponding LDAS data acquisition systems at the observatories. Another possible Critical System is the tape archive at Caltech, depending on whether a careful analysis would show that it is or is not possible for a malicious intruder to destroy significant amounts of data (or metadata).

Designation of a Critical System is made by the Director, Deputy Director, or Senior Computer Security Officer. Each Critical System should have a detailed security plan signed by the SCSO. This plan includes a definition of perimeter protections. The firewall rules for a Critical System are set to Default DENY, and the plan defines precisely, and justifies, what traffic is permitted in and out. The plan also details any gateway computers, precisely what authentication is required to move from the gateway into the Critical System, and lists those authorized to do so (or names an individual who keeps a authorization list). No one has authority to deviate from the plan without explicit written permission from the SCSO.

Non-critical Computing

As noted above, each autonomous computing unit is on one or more subnets with individualized firewall rules. For non-critical systems the rules are to start at Default PERMIT. Each restriction must have a reason and its implication on science and other computing evaluated. The plans must include explicit consideration of the operating system upgrades and patches. There is a trade-off

in the impediments caused by firewall restrictions and the effort and disruption to programs caused by extremely prompt implementation of patches and upgrades that can allow one to live with a looser firewall.

Putting firewall defenses at the perimeter of subnets corresponding to autonomous computing units rather than at site boundaries avoids a one size fits all approach and allows each unit to determine the level of restrictions consistent with the networking flexibility required within that unit to get its computing done.

This topology does come with one important complication. Many scientific users operate within a number of units and may often want to cross-mount files between units. Nothing fundamentally precludes this and the units can enable access through the firewalls to other units both at the local site and at other LIGO sites. Where the problem comes in is with regard to authentication and authorization of users at different units. Each unit maintains its own authorization lists, but if the units use different authentication mechanisms, the burden on the scientific users will likely be intolerable. Authentication is discussed below.

Network configuration and agility

The rearrangement of site networks into subnets corresponding to autonomous computing units will require some significant effort but, I am told, this is not overwhelming given the benefits.

The number of subnets likely to be required is within the range of existing Cisco routers. A large increase in the number of firewall rules could lead to increased CPU demands on the routers, which can be upgraded. This needs to be examined in a detailed network implementation plan.

A single network administrator at each site is the only one authorized to maintain and make changes to the rule lists on the routers. This might seem like a large job given requests coming from multiple computing units. However, Shannon Roddy tells me he has been using a software tool with a good interface that makes rule changing on routers relatively straight forward, far easier than the traditional line by line entry of rules.

Authentication

To be clear about definitions: authentication is the mechanism for establishing the identity of a user; authorization is permitting an authenticated user access to a system or a data file. Robust authentication is an essential protection mechanism for systems connected in any way to the external network. It should be easy for users or they will establish detours around it that will compromise essential security.

As noted earlier, different authentication mechanisms on different computing units would be a significant problem. It would be ideal if LIGO could converge on one authentication scheme for all systems. Users would be delighted to have to deal with only one password or a single one-time password (Smartcard) system.

Unfortunately different demands from operating systems, site histories, external sources (what will the DOE/NSF grid do?), etc. makes converging to one scheme something that will require considerable thought, negotiation, and (ahem) compromise. Here all sites (including the Tier 2 centers external to LIGO Lab) and all “autonomous” units would have to give up some

autonomy. So, the obvious solution at the moment is to establish a working group with representatives or knowledge from all entities.

Converging may not be as bad as it first looks. The DOE/NSF grid will have to deal with the different authentication schemes being used at major HEP labs, and LIGO is likely to move to a solution that is used at least at one of these labs. It would be useful to understand how SLAC, Fermilab, and CERN are thinking with regard to grid authentication to their sites.

Rapid response

I do not believe that LIGO needs to get carried away with an extensive rapid response team. However, there needs to be clear lines of authority during a security incident and a well defined, 7x24x365, contact point to report incidents. If an incident only affects one unit, that unit should clean up the mess (which will motivate a reevaluation of its security plan). The determination of whether an incident only affects one unit must be made at a higher level and so all incidents need to be reported centrally. Incidents on critical systems also require central high level management.

Rules and enforcement

Lots of rules get ignored. LIGO should have a short list of key computer security rules, including: No Hacking, Report Incidents, Don't Change Critical System or Network Configuration Unless Authorized, Install Critical Patches (as defined by the LIGO Senior Computer Security Officer), etc. Maybe 8-10 rules in all, half a page or less, and enforced just like similarly serious safety violations.

T. Nash 7/13/2004 DRAFT