Securing a HENP Computing Facility

Shigeki Misawa, Ofer Rind, Thomas Throwe
RHIC/US Atlas Tier 1 Computing Facility
Brookhaven National Laboratory
March 24, 2003

Facility Overview

- Supports data processing requirements for HEP/NP experiments.
  - Archive data store
  - Batch/Interactive Compute Farm
  - General interactive login facilities
- Also provides some resources for experiment administrative and support staff.
- Facility is embedded within a larger site.

Facility Structure (Before)

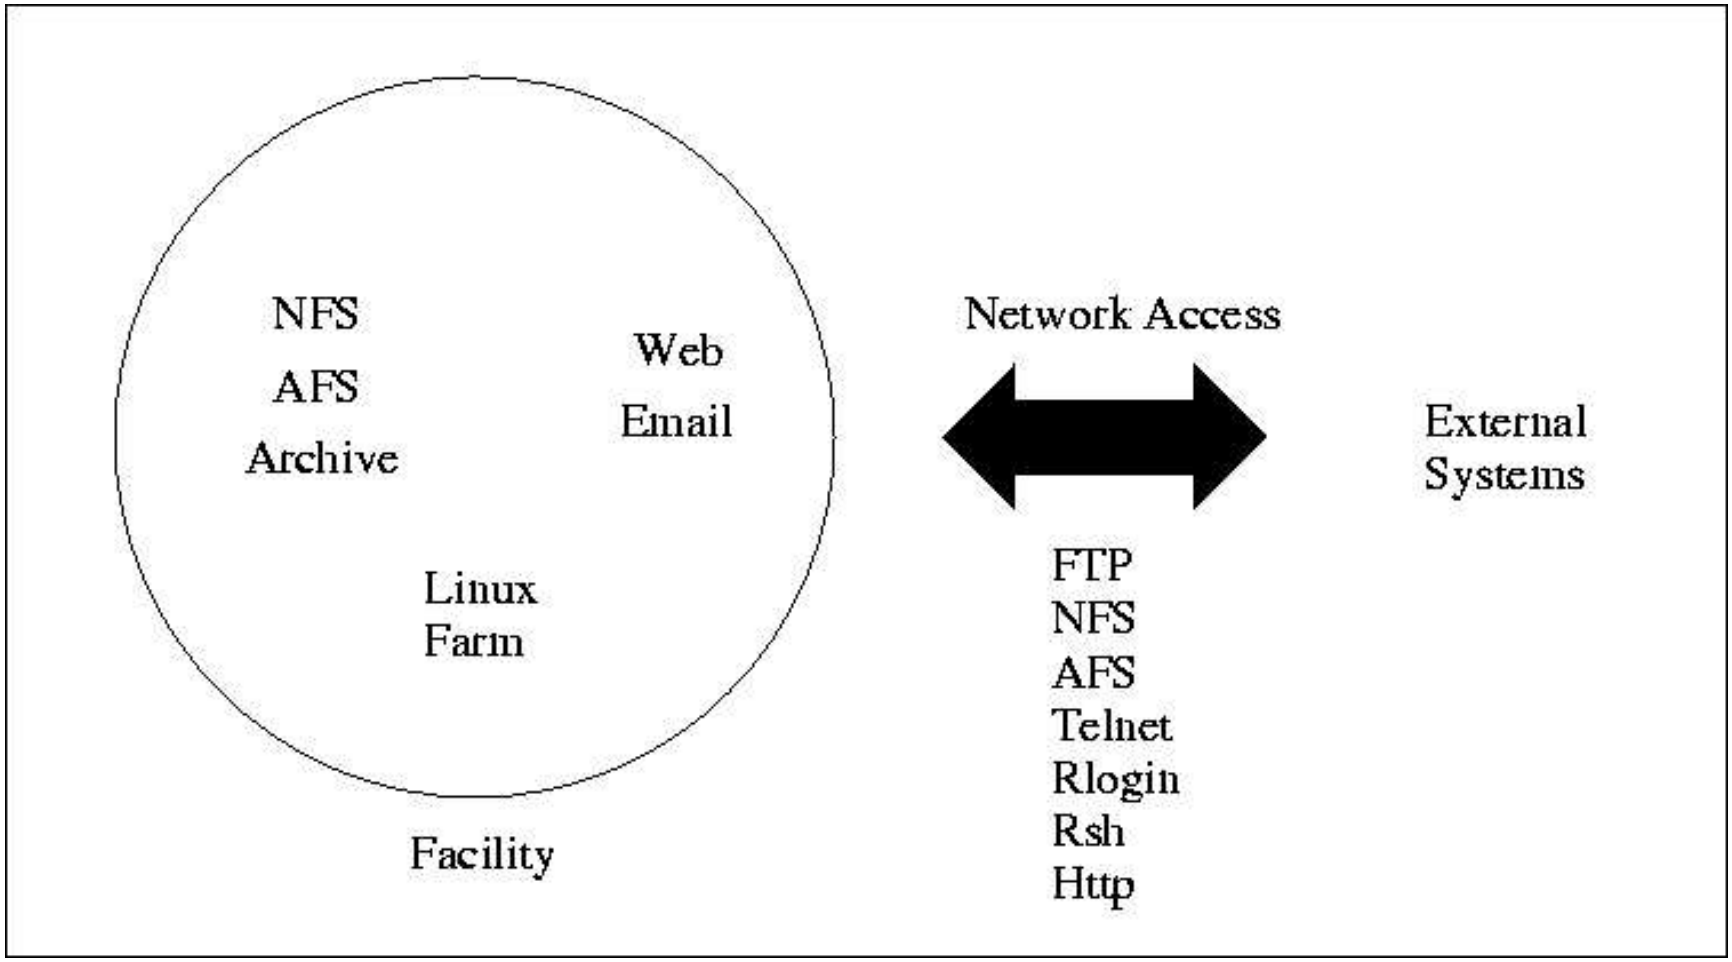

Securing a facility: Why ?

- Recovery costly in time and money.
  - Repository of hard to replace information.
  - Enormous facility “state” (10's of TB of data, 100's of systems).
- High Profile Target.
  - US Government Facility.
  - Vast resources, perfect (D)DoS launch site.
- Liability (if used to launch additional attacks.)

Identification of Assets

- Data in archive data store (HPSS)
- Central disk storage (NFS storage)
- Wide area accessible storage (AFS storage)
- Computational resources (Linux Farm)
- Distributed disk storage (Linux farm local disk)
- Misc. infrastructure services (KRB5, NIS, etc)
- Web services

Assessing the Threats

- Direct Network Assaults
  - Direct attack of network services
  - Hijack of network connections
  - DoS and DDoS
  - Web based attacks (server AND client side)
- Compromised Accounts
  - Stolen/cracked passwords
  - Social engineering
  - Dead accounts/malicious insiders

Responding to Threats

- System Architecture (The easy stuff.)
  - Network Topology
  - Technology Choices
- Human Factors (The hard stuff.)
  - Process and procedure
  - Awareness
  - Motivation

Why Prioritize ?

- Trade-off between security and ease of use.
- Resource hardening varies in difficulty.
- Economics – Hardening require time and effort.
- First step in compartmentalization of facility.
- Useful in the design of triage system in event of security breach (as well as “normal” system failure.)

Prioritization of Asset Protection (Decending order)

- Integrity of data in archive data store.
- Archiving of experiment raw data.
- Integrity/operation of the core facility.
- Interactive access to core facility from on-site
- Interactive access to core facility from off-site.
- On-site access to other core facility services.

Prioritization of Asset Protection

- Off-site access to other core facility services.
- Integrity/availability of email services.
- Integrity/availability of web based services.

Identification of Data Flow

- Counting house to Archives.
- Archives to/from NFS/Linux Farm Storage.
- NFS/AFS to/from Linux Farm.
- NFS/AFS to Web Servers.
- NFS/AFS/Linux Storage On and Off Site.
- Login On and Off site.
- Web/Email service On and Off site.

Designing the Facility

- Core facility in firewall protected zone.
  - Decrease observable cross-section.
- Multi-component firewall for scalability.
- Partitioning and zoning of web services.
  - CGI vs Static, Auth vs No Auth, Resources used.
- Decoupling of unrelated services (e.g. Email).
- Introduction of more secure services.
- Elimination of insecure services.

New Facility Configuration

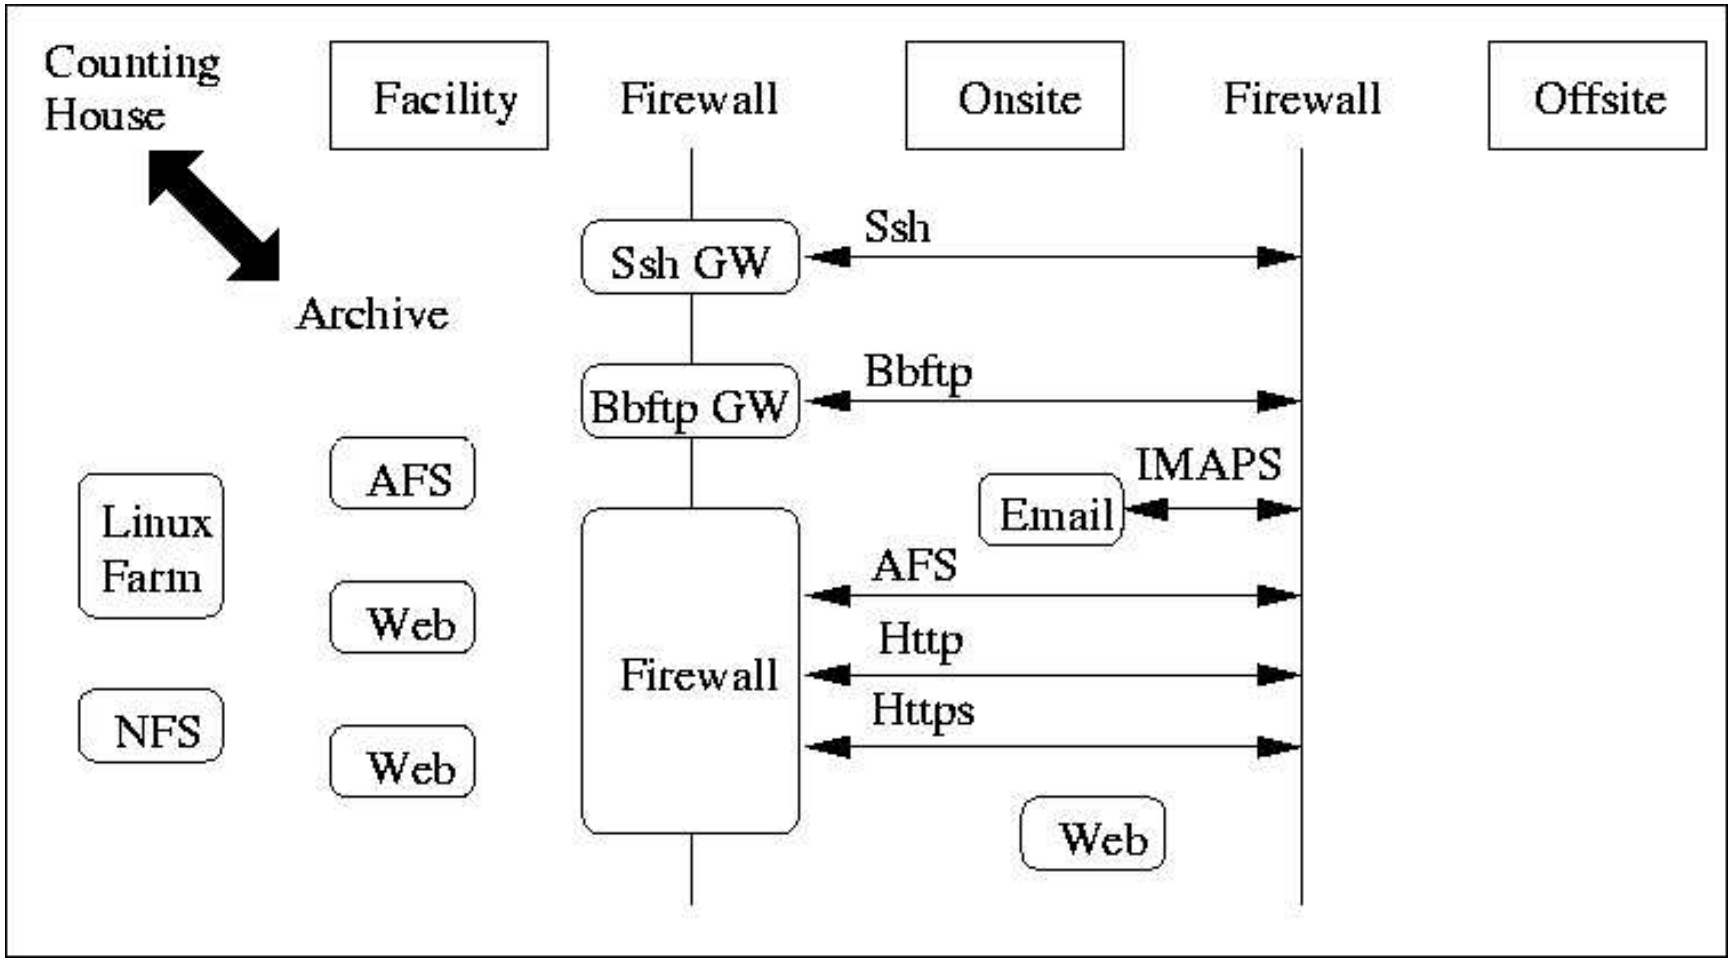

Core Facility
in protected network

- Layered core (logical)
  - Archive Data Store and counting house.
  - NFS, AFS and Linux Farm all in protected network.
  - Connect in -> out, not out -> in.
- Triage Scenarios:
  - Disconnect from Internet
  - Disconnect from campus network
  - Shutdown NFS/AFS/Linux Farm.

Facility Firewall

- Inbound “Default Deny” policy.
- Ssh “gateways”
  - Hardened (w/ipf), dual NIC systems with local home directories.
- FTP “gateways”
  - Hardened (w/ipchains), dual NIC systems for bulk data transfer purposes.
- “Normal” firewall appliance
  - Protect other network services.

Web Services

- Multiple Web servers.
- Non-authenticated CGI (CGI assumed to be evil)
  - Server outside of facility.
- Authenticated CGI (Assumes benign auth. User)
  - Server inside facility (requires facility services).
- Static user pages (No CGI)
  - Server inside facility (requires facility services).

Email Service

- Standalone mail server.
- SSL protected Web email access.
- SSL protected POP/IMAP email access.
- Outside of secure computing facility.
- Separate password database.

Technology Choices

- Ssh -> (replace telnet/rlogin/direct X access.)
  - Interactive access/X11 protection.
- Bbftp -> (replace ftp)
  - Data exchange with encrypted passwords.
- Kerberos -> (replace NIS passwords)
  - More secure authentication.
- SSL -> (encrypt clear text connections)
  - Secure authenticated web/email services.

Human Factors

- User/Administrator education
  - Can help but not a panacea
    - Awareness vs understanding.
    - Theory vs practice.
    - Documented procedures help.
  - Augmented with enforcement
    - Proactive password checking.
    - Encrypted connections.
    - Disabling of “bad” services.
    - Firewall rules prevent some bypass of security.

Operational Experience

- Ssh Gateways – working well
  - Celeron 650Mhz / 256Meg 120-140 users/system
  - CPU utilization < 30%
  - 1GB swap in use
  - Single graphics intensive app can saturate CPU
  - Scp transfers also an issue.
- Config/Maintenance/Upgrade an issue.

Operational Experience (cont'd)

- FTP Gateways – some issues
  - Sftp/scp (user) preferred transfer mode.
  - Bbftp not like ftp – barrier to widespead adoption.
  - Data transfer from distributed Linux disk is problematic. (Really an architecture issue.)
  - NFS disk appears to be a performance issue. (Faster central disk ? local “scratch” disk?)
  - Firewalling problematic.
- Config/Maintenance/Upgrade an issue.

Operational Experience (cont'd)

- Web services – working well.
  - Update mechanism is an issue.
  - Verifying integrity of content an issue.
- Firewall maintenance is an issue.
  - Procedures and process need development.
  - Gradually turning into leaky sieve.
- Intrusion/attack detection needs improvement.
- Authentication/Authorization management problematic.

Security Scorecard

- Facility is now significantly tighter.
- Weaknesses still exist.
  - Reusable passwords still weak.
  - Limited defense in depth.
  - Additional “conventional” hardening still possible.
- Security architecture is slowly being subverted by changes in operational requirements and facility expansion.

Scorecard

- Internet and threats have moved on.
- Defenses need to be upgraded to handle current threats (not the threats of the 1990's).

Future Work

- Tighter security on all deployed “edge” systems.
  - Better understanding/preparation. New technology?
- Tighter firewall configuration
  - More paranoid stance.
- Protection of web clients, not just servers ?
- Better pro-active user management ?
- Single Password/Single Sign on?

Grid Issues

- New types of security problems.
  - Globally accessible authentication/authorization services.
  - Stateful compute nodes -> sleeper programs.
  - Hijacked or infected executables/data.

Grid Issues (cont'd)

- New hurdles to overcome.
  - Distributed responsibilities. (Who does what ?)
  - Distributed authority. (Who's in charge?)
  - Distributed “facility”. (necessary resources everywhere)
  - Harder communication problems.
  - Additional complexity.
  - Non-mainstream services.

Grid Issues (cont'd)

- Phase in of grid services can be problematic.
  - Incomplete/untested services
- Adds new scale to ramifications of breaches in security.
  - Access to global computation resources
  - Destruction of global resources

Conclusions

- Security is hard.
- Security is an on going process.
- Low hanging fruit has been picked.
- Security is getting harder.