Secured Advanced Federated Environment (SAFE): A NASA Solution for Secure Cross-Organization Collaboration

Edward Chow PhD[1], Matthew Chew Spence[3], Barney Pell PhD[4], Helen Stewart[2], David Korsmeyer PhD[2], Joseph Liu PhD[1], Hsin-Ping Chang[1], Conan Viernes[1], Andre Goforth[2]

[1]*Jet Propulsion Laboratory*
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109

[2]*NASA Ames Research Center(ARC)*
Moffett Field, CA 94035
[3] *QSS Group Inc@ARC*
[4] *Research Institute for Advanced Computer Science(RIACS)@ARC*

This paper discusses the challenges and security issues inherent in building complex cross-organizational collaborative projects and software systems within NASA. By applying the design principles of compartmentalization, organizational hierarchy and inter-organizational federation, the Secured Advanced Federated Environment (SAFE) is laying the foundation for a collaborative virtual infrastructure for the NASA community. A key element of SAFE is the Micro Security Domain (MSD) concept, which balances the need to collaborate and the need to enforce enterprise and local security rules. With the SAFE approach, security is an integral component of enterprise software and network design, not an afterthought.

1. Introduction:

Like many federal agencies, National Aeronautics and Space Administration (NASA) field centers are distributed across the United States. NASA contractors and partners are located throughout the world. Many NASA projects and missions involve geographically distributed teams at NASA centers, industry, and universities. In 1999 the NASA Collaborative Engineering Environment (CEE) project developed a proof of concept "CEE room" to allow engineering teams from different sites to see and hear each other, share design data, and view and manipulate CAD drawings together in real-time over ISDN.

Prototypes were deployed in all ten major NASA installations within one year. One deep space mission required collaboration between a NASA center and a NASA contractor site. Using CEE rooms to collaborate saved the spacecraft design team 70 person-trips in the first two months, recouping the $70K equipment investment in 60 days. The CEE room concept was adopted as a standard by the NASA and continues to be used to support NASA mission design teams.

We began investigating extending these capabilities over IP to engineering desktops to enable projects to create teams, composed of members from different locations and employers, capable of sharing access to project-specific resources as well as CEE functionality.

We did not realize at the time that we were entering a minefield of organizational boundaries, inter and intra-enterprise relationships, differences in institutional priorities, and even some fundamental problems with the ways that network security is performed today. Our initial estimate of a 12 month project turned into a four year journey.

2. The NASA Environment

The nature of the American political system [1] encourages that NASA programs mete out development of complex government systems in a distributed manner across multiple states. Major NASA programs need to coordinate work across several NASA sites, NASA contractor facilities, associated university researchers, space-act agreement partners, other government agencies, and international space agencies. The competitive nature of some of the organizations involved can exacerbate the situation. The creation and operation of large government systems is generally contracted out to multiple vendors to perform together, and it is not uncommon for companies working together on one project to be in direct competition on several others.

NASA project and mission managers are frustrated with the loss of control over project timelines due to inter-site security issues. A fundamental problem is that the people who run IT security at a NASA Center have

widely divergent priorities and mandates from those of the people who run projects and missions.

While the project manager is charged with coordinating dispersed team members towards the project goals, the organizational security personnel, being responsible for the enterprise, site, division or group, are accountable for protecting the IT resources within their area of responsibility [2]. It is challenging to reconcile the distributed requirements of NASA-wide programs with the site-specific security procedures of multiple local domains.

Enabling researchers and engineers at different sites to share resources is a convoluted task, requiring approval of each researcher's organizational security and that of the domains in between. When sensitive data is involved, this can be extremely difficult between NASA facilities, and is aggravated further with contractor, academic, industry, and foreign partners.

Unlike many other governmental facilities, NASA Field Centers have a high level of autonomy, similar to that of universities within a large state higher education system. Each site has a distinct subculture, influenced by the primary center activities, the corporate cultures of on-site organizations, and other factors. Connecting researchers at three NASA sites via IP may require transiting nine or more organizational domains. Every organization has independent IT security processes, often with minimal coordination with other domains. IT security processes and mechanisms, although state of the art, tend to be domain-centric.

To avoid this complexity, projects create project-specific stovepipe networks and security domains. Other government agencies take a similar approach by building activity specific isolated networks[3], requiring installing and maintaining multiple infrastructures. Multiple computers sit on a user's physical desktop: one for Internet access and one for each type of secured access. Ironically, reducing the proliferation of keyboards and screens necessary to access separate systems was the impetus four decades ago for ARPA to build the ARPANET [4], which led to the modern Internet.

2.1 Cultural Impediments to Change

Organizational security teams may regard a cross-organizational system as antithetical to their mandate to protect their organizational perimeter, since inter-connecting with systems under external control may compromise the security of their organizational domain.

Project managers may view standardized approaches for inter-organizational collaborations as a risk to accomplishing specific project goals within a limited budget and schedule. Sharing information and resources across sites may not match preferred business methodologies or legacy project systems and may not contribute directly to the project goals.

NASA partners may not support inter-organizational collaborations. Partners such as contractors, universities, and other government agencies typically are engaged in multiple simultaneous cross-domain projects. Partnering with another organizational and security management may be perceived as too risky.

Finally, local NASA management may not support inter-organizational collaborations.. It takes vision and courage for management to support fundamental changes in key organizational business processes.

3. Defense in Depth

There are two principal ways to "firewall" systems to protect organizational security. Perimeter-based security, uses stateful IP firewalls as virtual portcullises to block unwanted traffic from entering a domain's virtual "castle walls". Host-based security places controls at each individual host, which requires configuration changes to underlying host operating systems and potentially installation of additional software as well. NASA sites, similar to many large organizations, use a multi-tiered combination of both to provide "defense in depth."

3.1 Perimeter Defense

Enterprise class firewalls are designed to protect many shared resources behind them. An enterprise class firewall needs to support gigabit per second throughput from multiple networks simultaneously while handling dynamic protocols such as H.323 voice & video and Web Services. Institutional firewalls require a team of trained professionals, and do not provide an absolute guarantee against compromise.

The IT resources of many projects may be cross-accessible on the institutional infrastructure behind the firewall. If a computer owned by one is compromised, others may be at risk. An intruder may use this computer as a bridgehead to attack other resources.

To prevent this from happening, organizational security personnel follow strict policies designed to protect the network domain, not individual projects. Such policies can create problems for projects with off-site team members, requiring exceptions to the rules to allow certain projects to work with external partners. Security staffs are wary of creating rule exceptions or "holes" in domain firewalls as tracking these can be difficult. Technical and management staff turnover can results in lost institutional memory. Domain security staffs are not always notified of a project shutdown, resulting in zombie firewall holes. Rule exceptions have been known to dominate a security staff's work in maintaining the firewall.

3.2 Virtual Private Networks (VPNs)

Network security technologies such as Virtual Private Networking (VPN) allows remote users to access

resources behind an institutional firewall by using an encrypted tunnel at the network (IPsec), transport (TLS/SSL) and/or application (SSH) layer(s) of the OSI model. VPN-based private extranets are extensively used to address cross-enterprise collaboration issues. The advent of inexpensive VPN hardware allows NASA projects to create stovepipe solutions at a significantly lower cost than in the past. However it is not unusual for the VPN mechanism employed by one partner to be denied to cross the firewall(s) of another partner.

Revocation of access across employers is a serious issue. For example: a team member who works at site A by nature of employment has VPN access into the network of site B. When this team member is fired, without an effective inter-organizational revocation mechanism, this person can still connect into site B and cause severe damage. It is difficult to decide which site is liable. Many organizations take a conservative approach that offers limited, if any, VPN collaboration capabilities with partners.

Enforcing host-based security can be difficult in telecommuting VPN environments, especially with laptops as the install, configuration, and patching of the user system may be outside the control of enterprise security personnel. The complexity or inconvenience of using host-based security software may result in the person who operates the computer misconfiguring or intentionally disabling it, defeating its purpose [5].

4. Impact of Emerging Technologies

Emerging networked technologies are stretching current institutional IT security approaches. For example, port forwarding protocols like stunnel and SSH may render transport-level firewalls obsolete, as the TCP/UDP port number of a stream will no longer be sufficient to determine appropriateness of incoming traffic. Web services, computational grids, and ubiquitous wireless connectivity pose further problems.

4.1 Web Services

Web Services offer a mechanism for organizations to quickly adapt their enterprise information systems to handle rapid change. However Web Services security protocols and methodologies are still quite immature. The HTTP tunneling and application firewall mechanism used by Web Services assumes that the user application is clean and blessed by the institutional security staff. If there is vulnerability in the application or a user installs unauthorized software, crackers may gain access, enabling them to threaten other enterprise resources.

4.2 Computational Grids

The GRID computing Virtual Organization (VO) has developed a good way of handling resource sharing across multiple domains [6]. Resource owners can determine how, when, and what is available for access [7]. This is a good model for scientific computing where scientists have complete control over their server and applications. In many organizations, resource owners might not have final say on resource availability. If International Traffic In Arms Regulations (ITAR) controlled data is involved, institutional security teams and the international affairs office will have to approve the access from outside of the organization. While a good start in assuming a homogeneous security environment, the GRID resource-owner-controlled model needs to be improved to accommodate the hierarchy of organizations involved within an institution.

4.3 Wireless

Technology is surpassing the traditional network firewall. Wireless network technology is becoming ubiquitous. Many new laptops have at least one wireless network interface in addition to wired Ethernet connectivity. It is not that difficult for a technically inclined individual to configure such a device as a router, allowing systems within wireless range access to the institutional network, possibly bypassing the firewall. 3-G wireless Internet access is already occurring in Japan and Europe. The traffic of these wireless networks is not going to go through the organization's firewall [8]. It will be difficult to enforce a policy to not use wireless devices as they will be pervasive.

4.4 Problem Summary

A summary of the security problems we faced when we started to build our collaborative environment are:

- Application traffic blocked by firewall.
- Conflicting organizational security priorities and project collaboration priorities.
- Confusing and uncoordinated processes for submitting multi-domain project security requests.
- Difficult distributed user authentication, authorization, and management.
- New technologies challenging the effectiveness of traditional firewall-based security.

5. SAFE Architecture

The SAFE project is a multi-center collaborative research project between the following NASA sites: Ames Research Center, Jet Propulsion Laboratory, and Marshall Space Flight Center. The SAFE project has created a unified architecture to give project managers the maximum flexibility to run their project and deploy their software while giving organizational security teams a unified way to manage collaboration security for multiple projects across multiple administrative domains.

5.1 *Goals*

Our primary goal is to prototype a collaborative environment within which NASA project teams can securely deploy software applications and share and exchange information and tools with their partners from around the world. We intend to develop for NASA a unified means to provide shared application services, and give NASA partners secure trusted teaming methods.

A collaborative environment is not effective if users cannot leverage upon the teaming mechanisms built by other projects. Hence another goal is to reduce new "stovepiped" project infrastructure implementations. The end result of our project is a multi-layered collaboration fabric linking knowledge resources, where multi-modal information, tools, and domain experts are available to NASA engineering teams when needed.

5.2 Architecture

The SAFE system includes a rack of equipment called a SAFE Node, which resides physically within a NASA center but logically outside the institutional network security perimeter. The SAFE virtual infrastructure consists of nodes at participating centers interconnected with sibling nodes via encrypted tunnels into a *SAFEspace*, a distributed secure network and application environment for virtual teaming.

SAFE is an overlay virtual fabric concept similar in implementation to the original MBone [9] and the 6Bone [10] Network tunnels are used to allow participants to overlay an advanced networking technology over arbitrary combinations of intermediary IP networks and autonomous security domains.

5.3 Micro Security Domain (MSD)

The SAFEspace can be subdivided into project-specific *Micro Security Domains* (MSDs) consisting of an isolated union of project resources across autonomous network boundaries. The level of abstraction provided by SAFE allows MSDs to use a location-independent project-centric security policy enforcement model rather than the current organizational perimeter firewall model.

A distributed project can form a MSD across enterprise boundaries. A subteam can form a MSD that is a subset of the larger, now hierarchical project MSD. The resources within a MSD can communicate internally but not external to the MSD except through specified logical gateways.

MSDs are enforced by Java-based *micro firewalls* downloaded onto a users computer upon entrance to SAFEspace External access to servers within the MSD is denied by default. The access to servers within the MSD is granted per application or by default if the policy allows. If the project needs to access services outside the MSD, it must come through an access gateway which is controlled by the organization security teams. Conceptually the whole enterprise is composed of many MSDs with a highly controlled gateway.

5.4 Virtual Relocation

One subsystem in the SAFE node serves as a gateway into SAFE for *virtually relocated* project-specific resources such as compute servers, 3D scanners, or other project or subteam resources to be shared across sites via an MSD. The project resource becomes mapped into the SAFE MSD for that project and thereby accessible to any other project member using SAFE network access. Because the project specific applications are well known by the project teams, it is easier to allow or deny access to users without the worry of impacting other projects. Virtual Relocation is intended for multi-user servers maintained by professional system administrators and under change management appropriate for security level.

5.5 Common Services MSD

A key benefit of SAFE is that NASA can put different types of agency-wide services into different *Common Services* MSDs. Depending on the security requirement, project managers can dynamically allow or deny access to common services. These Common Services MSDs can play an important role in long-term knowledge management for NASA by enabling shared repositories for document and knowledge management systems. The Proof of Concept SAFE Common Services MSD baseline is functionally equivalent to the desktop collaboration capabilities of the CEE rooms.

5.6 User Permissions Management

User connection to the SAFE is through a Java-based Secure Shell tunnel. A subsystem within SAFE nodes serves as a user gateway into a project MSD. The SAFE MSD access role definition is delegated to a project account manager. The project manager has to assign the user one or more functional roles, which govern which project resources are available to the user within the MSD. The policy for the firewall software is defined for the MSD by the project and it cannot be changed or disabled by the user without disabling SAFE access. The MSD application firewall *cloaks* resources, so the only services visible to a user are those to which the user has appropriate access.

5.7 Federation Proxy Server

Our *Federation Proxy Server* (FPS), referenced by the application firewall as part of user log-on, acts as an authentication proxy to allow users entrance by authenticating and authorizing against an externally managed Authentication and Authorization (A&A) system, such as a project or site managed user database. The FPS manages enforcement of rules within the

company, and serves as the access proxy to the remote services, if permissible, of the common MSD. The FPS enforces the rule set defined to ensure the security for MSD that crosses the organizational boundary. This can be intra- or inter-enterprise access. Later versions of the FPS will allow authentication against unions of independently managed A&A servers, so that site, project, employer, and resource owner(s) may all have to authorize before someone is given role-based access to defined capabilities upon a particular project MSD.

5.8 Micro Firewall

Upon successful user authentication against the FPS, the SAFE application firewall downloads a Java-based micro firewall onto the users' machine that limits user access to other resources when connected into a MSD. As an MSD only serves one project, the rules set on the micro firewall are simple. Many of us have consumer firewalls at home that are reasonably secure without requiring a trained technician to operate. Micro firewalls are simplified java based firewalls loaded into user workstations at MSD login.

A micro firewall is closely associated with the networking layer. It can be deployed in the same device as a router. The simplicity and close association with an individual network link gives the micro firewall the advantage of matching to a high speed network. Future network interface cards (NICs) could include ASIC Micro Firewalls.

A MSD consists contain multiple distributed micro firewalls working in coordination. Distributed firewalls work together under a shared rule set to protect the MSD. It may be connected with encrypted tunnels to other micro firewalls in the same MSD. The micro firewall gives the enterprise security team capability to monitor and control the activities within the MSD across enterprise boundary if policy allows.

5.9 Security Policy Editor

Another important characteristic of SAFE is automated hierarchical security. The SAFE project proposes improving the automation of the security approval process. We are developing a Security Rule Editor (SRE), a graphical user interface tool will allow an MSD administrator to input the security rules describing the project MSD business logic, allowing automatically configure the micro firewalls.

5.10 Ruleset Hierarchy

A necessary capability of the SRE is inheritance. The organization security teams can define security rules which are automatically imposed on the SRE that project teams use. Only when exceptions are needed will the project teams need to explicitly go through the security approval process. Otherwise the MSD creates an independent security domain, and as long as the project communication does not leave the MSD, the inherited rule set would be simple. Thus the locus of decision matches the locus of information within the hierarchy.

5.11 Federation between Hierarchies

Another important characteristic of SAFE is group-based federated user management. To support real-time collaboration across enterprise domains, a mechanism to manage users across enterprise boundaries is required. Inter-organizational trust issues surface almost immediately. How can we trust a user from a different enterprise if we do not know the person? We believe the answer is to look at how people work today. When a person requests access to a system, we typically verify with a known authority why the person needs access, what accessible resources, and the duration/limitations to such access. If we can automate this third party based trust process, where "A trusts B , B trusts C, so A trusts C," we may be able to achieve real-time dynamic security for federating users.

5.12 Architecture Summary

- SAFE nodes housing SAFE services are logically located outside site perimeter.
- NASA-wide common services available to SAFE users are housed within the SAFEspace (root security domain)
- Project specific-resources to be shared within a MSD are virtually relocated into the SAFE MSD via persistent encrypted tunnels across the network connections.
- An application firewall serves as user gateway into Micro Security Domains
- The Java-based micro firewall loaded onto user's system as part of SAFE login enforces MSD rules

6. Future Work

We are working on the FPS with integrated SRE, as part of a larger "Central Services" architecture. Once completed, NASA projects can dynamically create or destroy federated collaboration groups. The FPS will enable publishing and exchanging legally binding security policies and rules to a collaboration partner's corresponding system. Enterprise security teams can enforce security policy automatically within SAFE. We envision SAFE micro firewalls that can be designed into networking chips. These simple micro firewalls could be so cheap that they would be provided for every network link.

6.1 Service Autodiscovery

The SRE serves another function of creating a point of discovery of information and tools in SAFEspace. As project personnel need to enter data in a standard format through SRE, this information is available to the enterprise directory for discovery purpose. We are working on linking SRE definition format to standards such as UDDI to share the information within SAFE.

6.2 Chains of Authentication Instant Revocation

Security right revocation is critical. The rules defined in the SRE will be combined with the capability to read and verify from multiple disparate A&A directories in real-time. When a person is fired at Company A, they are removed from the company directory by the Human Resource Department. Although the person may still be in the a SAFE project authorization system, this person is denied access because the SRE rule requires a chain of authentication between the project authorization system and the company HR directory. This method enables instant unilateral revocation of access rights.

6.3 Group Identity Management systems

SAFE will need to cross-authenticate against multiple group-based trust mechanisms such as the Internet2 Shibboleth [11] project, Microsoft .Net Passport, and the Liberty Alliance [12] for projects involving academia, contractors, and other agencies. A cross-domain trust infrastructure such as SAFE needs to enforce group-based trust systems across multiple enterprises.

7. Conclusion

Typically, a security domain for any distributed project or software application is created on top of the existing network based on physical organizational boundary. Because of different requirements and the desire to have a defense in depth, an enterprise may build several network layers each with different levels of trust. The deployment of distributed projects becomes difficult and fragile as projects become tied to specific network architectures. Building trust across organizations is considerably more difficult than doing so internal to an organization.

In SAFE, security domains are project and application, not site, centric. A project MSD is a logical collocation of project and enterprise resources. SAFE decouples the security domain from the physical restriction. SAFE allows the deployment of NASA-wide projects upon a common connection framework that addresses the needs of the local as well as the global security issues. The SAFE technology is an important step toward functional based enterprise security. It allows an enterprise to achieve the balance between the need for inter-organizational collaboration and the need to protect enterprise security.

8. Acknowledgements

The Authors would like to gratefully thank the following people without whom this would have never happened: Shoshana Billik, Chris Buchanan, Scott Burchell, Mike Conroy, Pat Patterson, Conan Viernes, and Robert Wheeler.

This work was partially performed at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautic and Space Administration. Work was performed at the NASA Ames Research Center, and the Marshall Space Flight Center. Funding was obtained through the Computing, Networking, and Information Systems project within the Computing, Information and Communications Technologies Program, and the Engineering for Complex Systems Program.

9. References

[1] Tocqueville, Alexis de, *Democracy in America*, Colonial Press, New York (NY) 1899

[2] Oppenhiemer, D., D. Wagner and M. Crabb, *Systems Security: A Management Perspective*, USENIX, Berkeley 1996

[3] Bamford, James, *Body of Secrets: Anatomy of the Ultra-Secret National Security Agency*, Anchor Books, NY2002

[4] Hafner, Katie and Matthew Lyon, *When Wizards Stay Up Late: The Origins of the Internet*, Simon & Schuster, NY 1996.

[5] Yee, Ka-Ping, "User Interaction Design for Secure Systems," Proceedings of Fourth International Conferences on Information Communications Security, Singapore, 2002

[6] Foster, I, and C. Kesselman [eds], The Grid: Blueprint for a New Computing Infrastructure, Morgan Kaufman, San Francisco 1999.

[7] Foster, I,. C. Kessleman, G. Studik, etc "A Security Architecture for Computational Grids," Proceedings 5th ACM Conference on Computer and Communications Security, 1998.

[8] Internet Security Systems, Inc. X-Force Global Threat Operations Center, "Internet Risk Impact Summary, for January 1 to March 31, 2003" http://www.iss.net , April, 2003.

[9] Eriksson, H, "MBONE: The Multicast Backbone," Communications of the ACM, 37(8): 54-60, August 1994

[10] Rockell, R. and R. Fink, *RFC 2772: 6bone Backbone Routing Guidelines*, ITEF RFC Editor , 2000

[11] Erdos, Marlena and Cantor, Scott, "Shibboleth-Architecture Draft," shibboleth.internet2.edu, May, 2002.

[12] http://www.projectliberty.org