

Software-as-a-Service (SaaS) and Physical Security Management for Federal Systems

Adapting to the forces of HSPD 12, Convergence, and FISMA

April 18, 2008



Abstract

Working to meet the requirements of Homeland Security Presidential Directive 12 (HSPD 12)—a broad presidential mandate that requires all federal employees and contractors to use the new, secure Federal Information Processing Standard 201 (FIPS 201) identification card—has many government agencies trying to determine the most cost-effective way to implement a compliant Physical Access Control System (PACS). At the same time, many see the use of this new identity standard—and the infrastructure changes it will require—as an opportunity to also unify physical and logical access control (i.e., logging in to computer systems) as part of the same effort. Further complicating matters, those responsible for bringing their agencies into compliance are also trying to understand how the new FISMA regulations apply to the servers and control panels that make up these identity management and access control systems.

This paper provides a background on the relevant government mandates, initiatives, and regulation, going on to demonstrate how a commercial sector trend known as Software-as-a-Service (SaaS) has direct application to economically achieving the intertwined requirements of HSPD 12, FIPS 201, and FISMA—all while providing a set of ancillary benefits that are not available with any of the traditional approaches to physical and logical security or identity management.

Table of Contents

1. The Three Forces Shaping Federal Physical Security Systems	4
1.1. The First Force: HSPD 12/FIPS 201	5
FIPS 201 Requirements	5
FIPS 201 and PACS	6
1.2. The Second Force: Physical/Logical Convergence	8
What is logical access control?	8
Logical access credentials and SSO	9
Identity management—a primary tool for convergence	10
Why does identity management matter to federal agency security initiatives?	10
1.3. The Third Force: FISMA	12
Implications for Networked Computer Systems	13
Implications for Physical Security Systems	13
2. SaaS meets the Three Forces	15
2.1. Background: SaaS	15
2.2. Recent Growth in SaaS	17
2.3. Examples of SaaS	19
Personal Financial Services	19
Sales Force Management	19
Accounting	20
ERP	20
2.4. Multi-Tenant Software Architecture	20
2.5. SaaS in Security Management	22
2.6. Total Cost of Ownership	23
Tenant Suites	24
2.7. Robustness of SaaS	25
Availability	25
Disaster Recovery	26
2.8. Information Security	26

3. SaaS and the Three Forces	27
3.1. FIPS 201 and SaaS	28
Implement One Change, Not Many	28
Scale of SaaS Systems Reduces Redundant Enrollment	29
Wide Geographic Reach	29
One System, One Set of Updates	29
3.2. Convergence and Identity Management with SaaS	30
Streamlined Workflow	31
Reduced Interface Count for HSPD12 Ecosystem	32
3.3. FISMA with SaaS	33
Congressional Testimony on SaaS	33
Support for SaaS within the OMB	33
SaaS Addresses Many Requirements Off-the-Shelf	34
Reducing Cost of Compliance	34
Government Agencies Beginning to Use SaaS	35
4. Conclusions	36
4.1. Ancillary Benefits	36
Fewer Credential Databases to Maintain	36
Single Record of Personnel Movement	37
Lower Costs, Pay-As-You-Go	37
Anytime, Anywhere Management	37
Benefits for Tenants	37
4.2. Alternative Ownership Models for SaaS	37

1. The Three Forces Shaping Federal Physical Security Systems

Physical security systems in general—federal security systems, in particular—have become subject to numerous outside forces that were not a consideration for security managers even a few short years ago.

1

The most momentous of these considerations has been HSPD 12 and its corresponding technical embodiment, FIPS 201. These two federal government initiatives have set the security industry on a new course toward providing better credentials and more uniform management of them, while simultaneously precipitating a need for an almost wholesale replacement of many security products.

2

The convergence of logical and physical security—particularly in the form of using a common credential for both computer or network (logical) and physical access—has ushered in an era of unprecedented cooperation between the traditional IT community and the physical security community. All of this happening, mind you, at the same time that more and more physical security systems have begun to share the same IP networks with their logical counterparts, thus bringing these physical security systems under the purview of entrenched IT management hierarchies.

3

Lastly, reflecting a growing concern throughout society at large, FISMA applies higher information security and management standards to all computerized systems used within the Federal government. This now brings physical security systems under a whole new brand of scrutiny, and, along with it, much higher process hurdles for meeting new functional requirements and objectives (e.g., FIPS 201 compliance and physical/logical convergence) in a timely manner.

In this section of the paper, we examine each of these forces in greater detail in order to provide the necessary context for subsequently explaining how a growing commercial trend—Software-as-a-Service (SaaS)—provides the greatest leverage for adapting to these trends with the least delay and expense.

1.1. The First Force: HSPD 12/FIPS 201

HSPD 12 and its technical manifestation, FIPS 201, have provided the single-largest impetus in many years for the federal government to re-examine how it provides access control to the many assets under its authority. Together, these mandates have established the requirements for a common identification standard, one that pertains to credentials issued by federal departments and agencies to their employees and contractors for gaining physical access to federally controlled facilities and logical access to federally controlled information systems.

In accordance with HSPD 12, the FIPS 201 standard defines the technical requirements for an identity credential that:

- Is issued based on sound criteria for verifying an individual employee's identity;
- Is strongly resistant to identity fraud, tampering, counterfeiting, and terrorist exploitation;
- Can be rapidly authenticated electronically; and
- Is issued only by providers whose reliability has been established by an official accreditation process.¹

While many agencies will be faced with procedural challenges regarding the implementation of new standards for identity verification and card issuance, many agencies will also discover that their existing "legacy" access control equipment cannot process the new FIPS 201 credentials. Therefore, they have begun the search for upgrades and replacement products that enable them to meet the standards for physical access control.

The following sections highlight the aspects of the FIPS 201 standard that present challenges for legacy products in the physical security industry, and help to explain why these challenges are causing many decision-makers to look for fresh approaches to finding a solution.

FIPS 201 Requirements

FIPS 201 requirements are publicly available from NIST² and have been widely explained and analyzed by the Security Industry Association, the Smart Card Alliance, and various other security and IT publications.^{3,4} Federal government employees and contractors are now transitioning to FIPS 201 credentials, with several large card issuance contracts already in place across multiple agencies.

The FIPS 201 standard specifies the type of smart card credential (Personal Identity Verification, or PIV, card) that must be used, the processes for creating it, the cryptography used to protect data on the card, and the way that these smart cards communicate with the "readers" that are typically located near the doors of controlled access areas.

¹ Federal Information Processing Standards Publication: Personal Identity Verification (PIV) of Federal Employees and Contractors," National Institute of Standards and Technology Publication No. FIPS PUB 201-1 (Washington, DC: U.S. Department of Commerce, 2006).

² "FIPS Publications" (National Institute of Standards and Technology, December 28, 2007), <http://csrc.nist.gov/publications/PubsFIPS.html> (February 10, 2008).

³ "Physical Access Control System Migration Options for Using FIPS 201-1 Compliant Credentials," Smart Card Alliance Publication No. PAC-07002 (Smart Card Alliance Physical Access Council, 2007).

⁴ "Important FIPS 201 Deployment Considerations," CoreStreet (CoreStreet, 2006).

Because the standard is very explicit about the requirements on the cards and their corresponding readers, the General Services Administration (GSA) has been able to establish an Approved Products List⁵ that allows federal government buyers to know which card-related products are approved for HSPD 12/FIPS 201 applications.

FIPS 201 and PACS

What the FIPS 201 requirements do not address is what happens to all the information embedded on these new credentials after they have been extracted by a “reader” and electronically transmitted to a Physical Access Control System (PACS). The absence of any such derivative requirements for PACS architectures is an omission that has left the industry to identify those requirements on its own. Similarly, it has left the GSA without a basis for creating guidelines for recognizing and listing compliant PACS architectures.

With respect to PACS standards, the security industry and the government have attempted to fill the requirements gap through the formation of a variety of working groups, committees, and ad hoc studies—none of which are binding, but all of which are highly informative and point to several types of solutions for this problem, as well as transitional stages between the present and the desired end state.⁶

To understand what is at stake—that is, what sort of equipment and information systems are affected by these changes—the following diagram provides a generic view of a typical PACS architecture like those currently installed throughout thousands of government facilities worldwide.

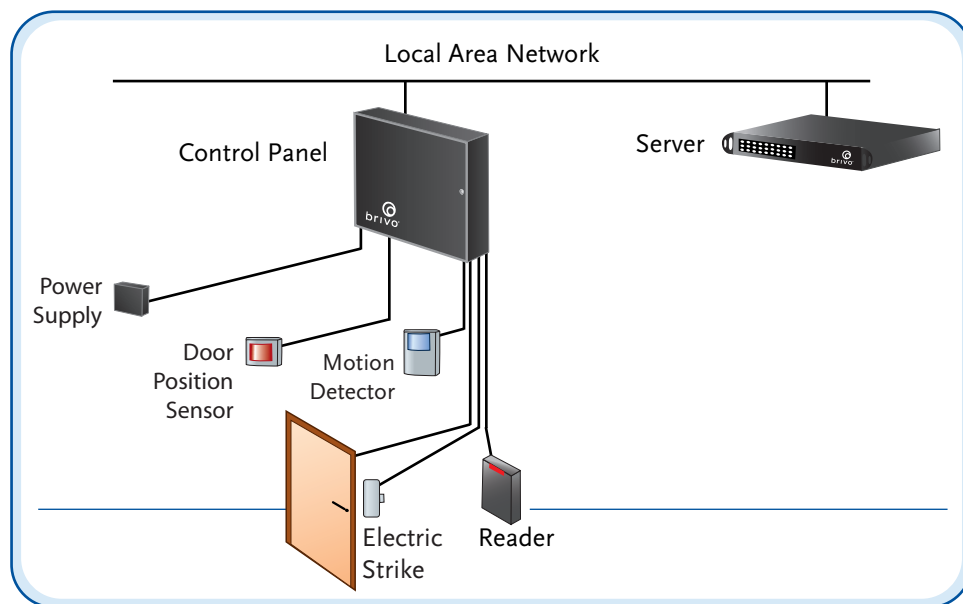


Figure 1: PACS Components

From these readings, it becomes obvious that one of the most salient aspects of the new FIPS 201 standard—from the perspective of a PACS—is that the amount of information (number of bits) in the new credentials is much larger than in traditional access control systems. So-called “legacy” systems—systems designed prior to the

⁵ “FIPS 201 Evaluation Program Approved Products List” (FIPS 201 Evaluation Program, March 10, 2006), <http://FIPS201ep.cio.gov/apl.php> (February 10, 2008).

⁶ The most definitive of these is the excellent Smart Card Alliance Publication No. PAC-07002, developed by the Smart Card Alliance Physical Access Council. See also the earlier “Considerations for the Migration of Existing Physical Access Control Systems to Achieve FIPS 201 Compatibility,” Publication No. PACS 06001, (Smart Card Alliance, September 2006).

FIPS 201 requirements—typically use credentials that range in size from 26 bits of information to 42 bits of information, for example.⁷

FIPS 201 credentials, by contrast, weigh in at a minimum of 75 bits of information, topping out at over two hundred. In a world where multi-megabyte files and email attachments have become commonplace, it is perhaps difficult to imagine that a mere handful—75 bits—of information is too big for a security system to accommodate. Yet that is exactly the situation that many physical security managers find themselves facing—with equipment replacements as the only solution offered by manufacturers.

The PACS equipment at issue in most of these discussions consists of two particular components: “control panels” and the servers (computers) from which the panels receive their instructions.

The control panel is connected to all the door-related electronics and is responsible for processing the credential data that a reader acquires when a card is “swiped” against it. Based on this information, the control panel either opens the doors or keeps it closed, and reports an invalid entry attempt back to a server somewhere. In order to meet the requirements of FIPS 201, this control panel equipment needs to have the necessary hardware and software capability to process the newer, larger credentials defined by the standard.

The second piece of equipment affected by this change is the server that’s running custom application software and manages one or more control panels connected to door-related hardware. These servers also need to interpret the new FIPS 201 card format. The challenge for the server software is not the size of the new credentials—after all, servers have more than enough storage—but rather the arrangement, or “format,” of the bits in the new credentials and knowing how to translate this data. The recognition of the new credential format must be programmed into the server in order for it to manage this new data correctly, associate it with users, and manage expiration dates and other information. Here again, most of the servers that have been deployed along with “legacy” PACS architectures do not understand these new credential types unless the manufacturer has updated them. If it so happens that the system in your building is no longer maintained by the manufacturer (i.e., it is out of date), you will be forced to upgrade to a newer software platform, which may, in turn, precipitate the need to upgrade the control panel hardware to hardware that is compatible with the newer server software.

Outdated and incompatible servers and control panels are the crux of the problems facing federal property and security managers today as a result of the HSPD 12 and FIPS 201 requirements. In the absence of an explicit standard for PACS, the security industry is left with describing its new PACS offerings as compatible with migration attempts to meet FIPS 201 requirements rather than compliant, mainly because there is no specific regulation with which to comply. At a bare minimum, though, it is clear that FIPS 201-compatible control panels⁸ and server technology need to be able to handle the information formats presented by the new standard.

⁷ While some larger credential sizes have been used in certain applications, the smaller sizes cited are by far the most prevalent in commercial real estate installations.

⁸ Throughout this paper, the term “control panel” is used interchangeably for all types of networked access control equipment that perform the same function, such as “single door controllers,” “edge devices,” etc. They are all logically equivalent and differ only in physical form and the number of doors they control.

Beyond that, there are additional functional needs of FIPS 201 systems that must be met by one (or more) components of the system architecture, including:

- Checking credential validity dates
- Checking credential status (e.g., revocation)

Examples of how these needs might be met are addressed in several of the references for this paper and will not be examined in detail at this point. We will, however, discuss how SaaS architecture makes it easier for these additional requirements to be consistently met across a very large population of individuals and properties.

1.2. The Second Force: Physical/Logical Convergence

Roughly coincident with the introduction of FIPS 201 and FISMA requirements, both the public and private sectors have gained a heightened awareness that standardized access control is not just an issue for physical assets, but also for “logical” assets—a shorthand term that refers to computer systems, software applications, file storage, or any other electronic representation of information.

In the words of the Open Security Exchange:

Today’s corporate security infrastructure is a patchwork. Most organizations maintain multiple, separate physical and IT security systems with no integration among them. This situation has become a growing liability as security concerns and the need to address privacy and regulatory compliance issues grow. At the same time, it prevents organizations from realizing an array of cost, control, and responsiveness benefits.⁹

This recognition is one of several factors that have given rise to a widespread focus on “convergence” between physical and logical security systems. In this context, convergence is defined as the merging of several operational and architectural aspects of physical and IT security systems to achieve a common management infrastructure—specifically with regard to combined credentials for physical and logical security.¹⁰

What is logical access control?

Most people encounter “logical access control” every day but don’t necessarily think of it in those terms. The phrase is simply the information industry’s name for the process of managing access to any sort of electronic information—a personal computer, a Web site, a corporate HR application, a financial database, and so on.

The most familiar form of logical access control is the use of a login name and password to get into a software application. Unfortunately, this simple form of access control has numerous shortcomings, ranging from the ease of maliciously appropriating someone else’s password to the frequency with which people forget their own.

⁹ “Physical/IT Convergence: What it Means, Why it’s Needed, and How to Get There,” Open Security Exchange (Open Security Exchange, 2007).

¹⁰ Note that the term “convergence” is widely used (and misused) in the physical security industry to refer to many different aspects of combining physical and IT security. For example, it is often used to mean nothing more than “IP convergence,” which refers to the sharing of a common network infrastructure to support both IT and physical security data traffic. While important from an installation cost perspective, the mere sharing of a common IP infrastructure does little to achieve any of the higher operational goals at stake for today’s enterprise.

An even bigger problem for enterprises with a large number of computer systems is that historically, each one has had its own individual repository of logins and passwords, or, more generally, credentials. This presents a major problem for managing users across multiple systems.¹¹ It is not uncommon for a new employee to have credentials—or an identity—established in at least four or five disparate systems upon joining a new employer (e.g., local area networks, remote access VPNs, email, file servers, physical security, specialized applications, etc.). And, of course, when the employee leaves the company, all those instances of credentials must be removed in order to protect data or remove access permissions.

Logical access credentials and SSO

The first volley at addressing this proliferation of logical access credentials has been to implement Single Sign-On (SSO) systems to manage all of a person's credentials in one place. This centralized approach makes it much easier for IT managers to add and remove employees, or otherwise manage access privileges to computer systems and data. More sophisticated logical access or SSO implementations also augment the relatively insecure login/password combination with more secure technology, such as biometrics, tokens, rolling codes, etc. Thus, IT systems have begun to mimic what physical access control has been doing for years: insisting on one or more factors (e.g., a smart card plus a PIN code, or even a fingerprint) to authorize access to computer systems.

That said, the organizational reality is that those responsible for physical security are most often a different group of people than those who are responsible for logical security. The result of this functional division is that each group has historically used different credentials to control access to their respective domains of responsibility—even if the underlying technologies were relatively similar. The effect is that while employees already carry a card or credential for physical access control, now they must also maintain one or more additional credentials for each computer system they need to access.

To again cite the OSE on this subject:

[T]his lack of integration is no longer simply an inconvenience. It increases security risks by preventing technologies from working in concert with one another. It limits corporations' efforts to establish centralized control of security and develop integrated risk management strategies.¹²

The question at hand then becomes: How do organizations get both types of systems—physical and logical—to use the same credentials? What are the steps that must be taken to achieve this goal? How does an organization minimize the expense of getting there?

¹¹ It is also very costly, with estimates ranging from \$400 to \$1,000 per year to provision and manage identities for workers. Jorgenson, Barbara. 2008. Identity Secured, Vol 6, Issue 3.

¹² "Physical/IT Convergence: What it Means, Why it's Needed, and How to Get There," op. cit.

Identity management—a primary tool for convergence

The ideal solution for all parties involved is to bring physical and logical systems into alignment so that they can all use the same credentials. This is one of the primary goals of convergence, with identity management being a key tool required to achieve that objective.

Identity management systems have been available in the IT world for a number of years. They are available as both pure software solutions and as hardware-based appliances. They are available as integrated suites, or as standalone solutions. And they are available as proprietary offerings that work with only certain operating systems, or as more open systems that work across a greater number of platforms. Regarding this last point, there are in fact a great many proposed standards for identity management, which is just one of the challenges facing convergence at this time. A recent survey of proposed standards for identity management lists no fewer than ten current candidates, including Shibboleth, Liberty Alliance Federation Framework, Security Assertion Markup Language, Web Service Secure Exchange, Web Services Federal Language, OpenID, Light-Weight Identity, CardSpace, Project Higgins, and Digital Identity Exchange.¹³

To date, most identity management systems have been applied after the fact; that is, after there was already a problem with a proliferation of multiple identities and credentials across many IT systems. What this spells for IT managers is a problem that becomes larger and more difficult to solve as the number of disparate systems increases. This scaling problem is a result of the fact that each application or access control system tends to have its own security model that may not easily map to those around it.

With more awareness focused on these issues, the expectation is that there will be better consideration given to identity management in advance of installing new systems, whether physical or logical. Among other factors, planning in advance for a centralized, authoritative credential database is key to reducing complexity and implementation costs.

Why does identity management matter to federal agency security initiatives?

Identity management matters to federal agencies because they need to ensure that new FIPS 201 compatible PACS architectures will be consistent with their own organizations' convergence and identity management goals. As a result, they need to know how to select vendors that can easily integrate with both logical and physical access control systems.

Because identity management is an arena where standards are still emerging, it is not possible at this juncture to spell out exactly what the characteristics of the ideal PACS solution would be. There will undoubtedly be a fair amount of volatility, with the accompanying need for flexibility in whatever solution is chosen. That said, the following guidelines provide a set of minimum requirements that any PACS solution should meet:

A recent survey of proposed standards for identity management lists no fewer than ten current candidates.

Proliferating identities is a problem that becomes larger and more difficult to solve as the number of disparate systems increases.

¹³ Jorgenson, op. cit.

PACS Requirement	Explanation
Uses IP Protocol	Using IP for communications among the major components of the PACS is an absolute requirement, and happily one that most current PACS architectures are able to meet. However, IP alone is no guarantee of any meaningful degree of interoperability with an identity management system (IDMS) in the absence of open APIs.
Open APIs	The PACS architecture should have open APIs that are accessible from any computer language, operating system, or hardware platform. Examples include XML-RPC, SOAP or any other variants of Web Services based interface. By contrast, APIs that require vendor-specific SDKs that only operate on, say, .NET platforms, are closed and will ultimately lock in the buyer to a proprietary architecture.
XML Data Interchange	While not an absolute guarantee of interoperability, the use of XML formats for data interchange are a strong indicator of the potential for easy integration with other systems.
Non-proprietary	Non-proprietary can have several meanings. First, it means interfaces that are not tied to one system architecture, as mentioned above regarding open APIs. More importantly, it means avoiding solutions that are heavily tied to the PACS architecture itself. ID management is an IT initiative that spans many application types with diverse requirements. PACS architectures are just one of many such applications, and they should not drive the entire IDMS solution. The IDMS solution should preferably be driven by broader IT considerations, while ensuring that the PACS system is compatible with those broad requirements, not vice-versa. Open APIs and XML data interchange are good steps toward accomplishing such compatibility.
Scalable	The size of potential PACS implementations in the federal government reaches numbers upward of millions of users and tens of thousands of door and monitored points. The more scalable the PACS solution, the easier it will be to administer. As discussed in the context of SaaS solutions, it is now possible to find systems that manage databases of this size, all in a single multi-tenant architecture.

Table 1, PACS Characteristics for Identity Management

1.3. The Third Force: FISMA

The Federal Information Security Management Act (FISMA) was enacted in December of 2002 in an attempt to strengthen computer and network security within the Federal Government and its contractors by requiring yearly audits. FISMA calls for each agency to “develop, document, and implement an agency-wide program to provide information security for the information and information systems that support the operations and assets of the agency.”¹⁴

As regulation that is applied to all information systems within the federal government, it will necessarily influence the way that any electronic security system is designed, managed, and monitored. As most observers agree, while FISMA is intended to boost overall security of federal information systems, there is little dispute that it will only add to the cost of any information system under its jurisdiction. What that means for access control systems is that not only must they be designed and installed in a manner consistent with FISMA, but they must do so with the least additional expense in the context of the “unfunded mandate” of HSPD 12 and FIPS 201.

FISMA requires specific processes to ensure the integrity of all federal information systems, with an effective program for information security to exhibit the following characteristics:

- Risk assessments regarding the degree to which information systems and related parties could be harmed by “unauthorized access, use, disclosure, disruption, modification, or destruction of information and information systems that support the operations and assets of the organization.”¹⁵
- Policies and procedures designed with the intention of evaluating security risks, the cost-effectiveness of these programs and their ability to reduce risk to a suitable level, and the assurance that the integrity is considered throughout the development of each organization’s information system.
- Auxiliary plans for providing sufficient information security for federal agencies.
- Training of information systems users as to the risks when using the system and their individual responsibilities to reduce these risks.
- Regular, if not annual, testing and evaluation of the integrity of the information system.
- “A process for planning, implementing, evaluating, and documenting remedial actions to address any deficiencies in the information security policies, procedures, and practices of the organization.”¹⁶
- Mechanisms and procedures for identifying, reporting, and addressing security events.
- Strategies to ensure the efficiency and continuity of operations of information systems.

Divided into two phases, FISMA attempts to improve the overall security of information systems through setting standards and guidelines, as well as creating a credentialing program. Phase One, completed as of 2007, defines information security standards, as well as identifying a guideline to assess potential risks to system integrity. Phase One also includes guides on how to properly certify federal information systems, risk management framework, assessing security controls, identifying an information system as a national security system, and information mapping.

¹⁴ FISMA Overview, (National Institute of Standards and Technology, February 15, 2008), <http://csrc.nist.gov/groups/SMA/fisma/overview.html> (February 20, 2008).

¹⁵ FISMA, op. cit.

¹⁶ FISMA, op. cit.

Implications for Networked Computer Systems

There are many aspects of FISMA compliance that IT managers and system users need to concern themselves with for each information system under their control. Many of the controls required are specific to certain types of systems, such as a Web server vs. a database server. Some requirements, however, are very general in nature, such as those listed in FIPS 200¹⁷, NIST's guidance regarding the selection of security controls for each system. Because of the generality of some of these controls, federal property managers can be certain that at least these minimum criteria will enter into accreditation of relevant electronic security infrastructure. For example:

- Media protection requires limiting access to media-related information on information systems to authorized users, and sanitizing or destroying such related information on information systems before disposal or release for reuse.
- Physical and environmental protection requires limiting physical access to information systems, equipment, and the respective operating environments to authorized individuals; protecting the physical plant and supporting the infrastructure for information systems; providing supporting utilities for information systems; and providing appropriate environmental controls in facilities containing information systems.
- Incident response requires establishing an operational capability for handling preparation, detection, analysis, containment, recovery, and user response activities, with the ability to track, document, and report incidents to appropriate organizational officials and/or authorities.
- Continuous monitoring requires ongoing determination of risks and compliance with specific guidelines outlined in the various special publications related to FISMA (800-26, 800-30, 800-37, 800-51, 800-53, 800-60).

At the end of the day, each installation of an information system anywhere in an agency needs to go through an approval process for that installation to make sure the system was installed in compliance with the relevant criteria. What this means is that the more systems that are installed, the more acceptance processes an agency must go through.

Implications for Physical Security Systems

Prior to FISMA, requirements like those listed above were not typically considered when installing PACS solutions. Standalone "legacy" systems were typically installed on a building-by-building basis, each with its own separate PACS architecture, its own user database, and its own IT architecture.¹⁸ Many such systems were not even connected to a network, which made IT risk assessment simpler¹⁹, but also made such systems very inconvenient to use and manage because they could not be remotely accessed, as shown in Figure 2, Legacy PACS Architecture – One per Site.

¹⁷ "Federal Information Processing Standards Publication: Minimum Requirements for Federal Information Systems," National Institute of Standards and Technology Publication No. FIPS PUB 200 (Washington, DC: U.S. Department of Commerce, 2006).

¹⁸ While some agencies have had the foresight to centralize certain parts of the security operations, separate and duplicate systems have been the dominant trend in commercial property management.

¹⁹ According to NIST Special Publication 800-53, "For example, when information system components are single-user, not networked, or only locally networked, one or more of these characteristics may provide appropriate rationale for not applying selected controls to that component." "Guide for Assessing the Security Controls in Federal Information Systems," National Institute of Standards and Technology Publication no. 800-53A (Washington, DC: U.S. Department of Commerce, 2007).

²⁰ "Guide for the Security the Security Certification and Accreditation of Federal Information Systems," National Institute of Standards and Technology Publication no. 800-37 (Washington, DC: U.S. Department of Commerce, 2004).

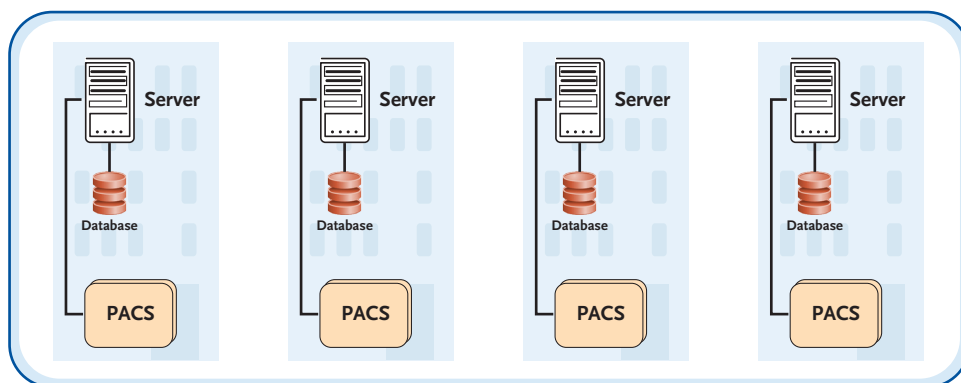


Figure 2: Legacy PACS Architecture-One per Site

With the legacy access control architectures, even systems that were connected to a network for such purposes as file backup or report printing were often not able to be remotely administered. However, because they were connected to a network, they did acquire all of the risk exposure that network connectivity implies, especially for certain vulnerable operating systems.

What this means today for new PACS solutions installed on a building-by-building basis is that at least certain aspects of each of them would likely have to be individually evaluated under FISMA Certification and Accreditation guidelines to ensure compliance.²⁰ The reason for this is that each one of them is potentially on a unique network with unique vulnerabilities, depending on how an individual property is connected to the Internet (e.g., ISP, firewalls, proxies, wireless, etc).

This implication of the FISMA C&A process effectively establishes a higher hurdle for each of these PACS solutions, and points toward the cost and operational benefits of shared, multi-tenant systems wherever possible.

²⁰ "Guide for the Security the Security Certification and Accreditation of Federal Information Systems," National Institute of Standards and Technology Publication no. 800-37 (Washington, DC: U.S. Department of Commerce, 2004).

2. SaaS meets the Three Forces

SaaS has emerged over the past five years as one of the strongest new trends in the computing industry. According to Nicholas Carr, the SaaS industry has been growing at over 20% per year²¹ and now represents a market size of approximately \$8-10 billion annually in the U.S. alone. Gartner similarly predicts that the SaaS market will continue to grow at 22.1% per year²² and that by 2011, 25% of new software systems will be delivered as SaaS applications.²³ Among the reasons usually cited for this steep adoption curve, primary motivators often include the ease of implementation and the ability to keep pace with rapidly changing requirements.

In studying the Three Forces, it quickly becomes obvious that the information systems required to perform the PACS function—when paired with FIPS 201 and convergence, and managed under the umbrella of FISMA—require considerably more sophistication than the systems they will replace. They are also likely to require frequent change, particularly in these early years of the Three Forces, as requirements in all three areas continue to mature. With greater sophistication and rapid change there is usually an increase in IT expense, a higher degree of difficulty to implement and manage, and greater concern for system reliability. The question is: How best to manage all this change?

For the federal property or security manager, adapting to the Three Forces will require systemic changes to the IT infrastructure used for managing access control, identity, and related physical security functions such as video surveillance. Opportunistically, this process of adaptation will also provide a unique venue for establishing uniform security systems architectures across an entire agency. Achieving such a consistent infrastructure at the lowest organizational and financial cost, while keeping up with mandated implementation deadlines, is the subject of the next section of the paper.

Specifically, we look at how the growing trend of SaaS delivers a ready answer for achieving these objectives. As recent studies indicate,²⁴ SaaS is extremely effective at lowering Total Cost of Ownership (TCO) for physical access control, which makes it a very desirable model in this instance.

2.1. Background: SaaS

SaaS is a software application delivery model where a systems provider develops a Web-native software application and hosts and operates (either independently or through a third-party) the application so that its customers can use it via the Internet. Customers do not pay for owning the software itself; rather, they pay a small fee to use it. Customers use the applications through either a browser or, programmatically, through an API accessible over the Web and often written using XML.

Gartner predicts that the SaaS market will continue to grow at 22.1% per year and that by 2011, 25% of new software application will be delivered as SaaS applications.

²¹ Carr, Nicholas. *The Big Switch: Rewiring the World, from Edison to Google*. New York: Norton & Company, Inc., 2008.

²² Scheier, Robert L. August 20, 2007. "Your Data's Less Safe Today than Two Years Ago," *InfoWorld*, http://www.infoworld.com/article/07/08/20/data-is-less-safe_1.html (January 4, 2008).

²³ "Gartner: SaaS Market Heats Up." September 28, 2006 ebizq, <http://www.ebizq.net/news/7314.html> (January 20, 2008).

²⁴ "SaaS—TCO," Brivo Systems (Washington, DC: Brivo Systems, 2008).

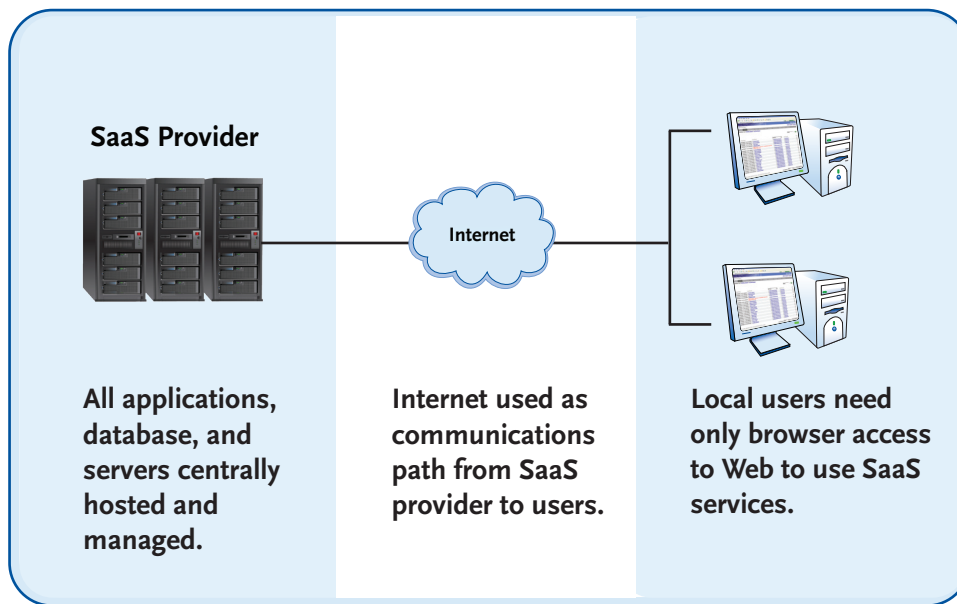


Figure 3: Generic SaaS Architecture

SaaS is at once both the most important new trend in computing and a return to its roots. Enabled by the ubiquitous bandwidth provided by the Internet, SaaS architectures transform computing resources from being localized in one physical “box” (like a PC) to being distributed as a network utility—literally in the same way that electricity and other utilities are provided on a grid and distributed to wherever they are needed. This parallels the way that computing resources were originally provided—by centralized mainframes. The big difference between then and now, however, is that the technology available today has made it possible for these centralized services to be provided as an on demand service, just like water or electricity—all without forcing the user to purchase, install, or maintain any local servers or applications.

For these reasons, SaaS provides huge efficiencies and economies of scale. First, it offers immediate savings on server hardware. A centralized SaaS computing model typically allows hundreds or thousands of distributed servers or PCs operating at low utilization to be replaced with a single server that operates at high utilization, thereby eliminating the expense of all those fallow computing resources. Second, all of the redundant electricity, cooling, and staff maintenance expenses disappear along with those underutilized servers and PCs. As even Microsoft points out, “The initial purchase is usually only 5% of the total cost of owning and maintaining a program.”²⁵ SaaS removes the majority of this expense from the equation by freeing the end user from having to own software and servers at every location where they are used.

“The power bill to run a computer over its lifetime will surpass the cost of buying the machine in the first place.”

— **Scientific American, April, 2008.**

²⁵ “Microsoft Wages Campaign Against Using Free Software,” The Wall Street Journal, December 9, 2002.

Many of these issues are discussed at length in Nicholas Carr's recently published book, *The Big Switch*, where he concludes that:

"The complexity and inefficiency of the client-server model have fed on themselves over the last quarter century. As companies continue to add more applications, they have to expand their data centers, install new machines, reprogram old ones, and hire even larger numbers of technicians to keep everything running. When you also take into account that businesses have to buy backup equipment in case a server or storage systems fails, you realize that, as studies indicate, most of the many trillions of dollars that companies have invested in information technology have gone to waste."²⁶

Just as SaaS has begun to eliminate much of this waste for private enterprise, it also offers a model for cost-effectively implementing security management systems within the federal context: SaaS offers its centralized management function as a shared resource to be used across an entire agency.

2.2. Recent Growth in SaaS

SaaS is one of the fastest growing segments of the IT industry because it provides a more cost-effective alternative for enterprises to achieve their business objectives than traditional packaged applications. It also provides greater flexibility for end users because they do not have to commit to the features and capabilities of a solution at just one point in time. Instead, they benefit from the continual evolution of the underlying software as new innovations become available.

Brivo Systems, for example, has customers who have been using its SaaS-based physical access control system for over six years, during which time they have automatically received no fewer than 25 upgrades to their core capability set—all without any of their own IT involvement or any additional expenses. These customers simply receive the latest version of the company's SaaS offering the next time they log into the system after an update (much like the additional features a bank customer receives from time to time when his or her bank updates its online service offering). Over time, this has protected Brivo's customers from obsolescence or costly software and server upgrade cycles, while still allowing them to enjoy the company's latest feature set on their existing control panel hardware.

In the broader IT market, SaaS has found many adherents for many different types of application services. The reasons for adoption vary by industry and application type, but a number of common trends emerge as reasons for adopting the SaaS model, among them being ease of deployment, flexibility, lower costs, and ease of use, as shown in the results of a recent study published by *Information Week*:²⁷

²⁶ Carr, Nicholas. *The Big Switch: Rewiring the World, from Edison to Google*. New York: Norton & Company, Inc., 2008: 56.

²⁷ *Information Week*, CMP Publications: April 2007

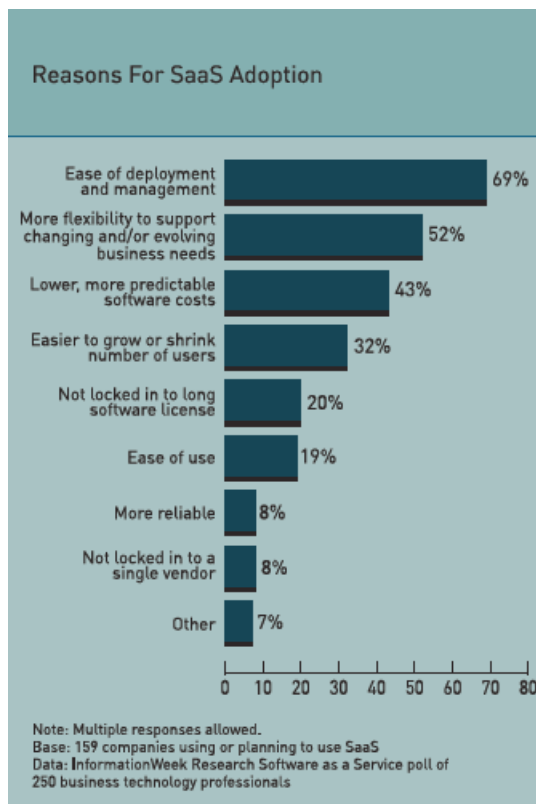


Figure 4, Reasons for SaaS Adoption. Source: Information Week.

2.3. Examples of SaaS

Many people are using SaaS already without even being aware that they are doing so. Many of us make use of the SaaS model in our personal lives, and, increasingly, many businesses are doing the same—even with the most critical and sensitive data their corporation possesses. Below we briefly examine several examples of such businesses, and then in more detail at how this model operates in the security area.

Personal Financial Services

Personal finance has been silently introducing all of us to the SaaS model since Internet use spread to most households in the U.S. For example, every time someone uses a bank's online bill payment service through a browser, they are using SaaS. Why? Because all of the software used to provide this service resides on systems outside of the laptop or PC where the user initiates these transactions. In this case, the user did not have to buy a special server, install any special software, or make any other specific infrastructure investment in order to make use of the bill payment service. The same is true for anyone who trades stocks online; the brokerage service does not require its customers to buy any particular software or servers in order to make these financial transactions. All a customer needs is a computer connected to the Internet.

Contrast this with the non-SaaS model that was prevalent just a few years before. In order to gain access to online bill payment services, for example, one needed to purchase, install, and maintain a product such as Intuit's Quicken or Microsoft's Money. These products were not only expensive, but required frequent upgrades and often became obsolete or unsupported over time, or when one's bank changed hands. Similarly, electronic stock trading was previously limited to those who were willing to purchase special software and/or direct connections to electronic quote services.

Sales Force Management

Perhaps the most-cited success story in the SaaS arena is Salesforce.com, an online service that allows companies to manage their sales force activities and compensation. The company was founded by a former Oracle Corporation executive, Mark Benioff, who realized that companies were not interested in managing their own software so much as they were interested in managing their sales:

While companies typically had organized their relationships with customers using customized software installed on computers, Mr. Benioff's idea was to offer a no-frills online service that let companies make only minor tweaks. Such Web-based software required less investment in on-premise software and hardware, got up and running more quickly, and was easier to use.²⁸

The most interesting aspect of the Salesforce.com success is that companies now trust the reliability and security of such SaaS systems as much as (or more) than they do their own internal IT systems. A company's sales database is absolutely critical to its success, and its customer database is among its most prized—and proprietary—information resources. To trust it to an outside organization speaks volumes about the level of trust now placed in the SaaS model.

²⁸ "Web Based Software Services Take Hold." May 15, 2007, SalesTechnology.com, <http://www.salestechnology.com/current-news/2007/5/15/web-based-software-services-take-hold.html> (January 15, 2008).

Accounting

Traditionally, small and mid-size companies that needed financial accounting software were forced in either one of two directions: inexpensive PC-based products that were feature-poor and didn't meet their needs, or very expensive server-based products that provided all the features they needed, but only at a very high cost, both in terms of initial investment as well as ongoing maintenance and upgrades. Into this mix came a new company, NetSuite Inc., a company that brought the SaaS model to bear on this problem. What the company provides is an Internet-based service that allows companies to manage not only their traditional accounting functions online, but also a wide range of enterprise resource planning-type (ERP) domains such as HR, inventory management, sales order processing, even e-commerce, all in a browser-accessible solution that requires no customer investment in servers or software licenses. NetSuite Inc.'s SaaS model has grown exponentially in the five years since its introduction, and the company recently went public with a successful IPO.

More noteworthy than the success of the company itself, however, is the fact that it, too, illustrates the level of trust that the market is now placing in the SaaS model. Obviously, employee and financial data are among the most sensitive and confidential information resources a company possesses. To be willing to entrust this data to an online service provider illustrates the degree of confidence that the public is now placing in this model.

ERP

Like the Salesforce.com example, the use of a SaaS provider for ERP software provides a powerful testimonial to just how far this new software model has come in the past eight years. In 2006, the German ERP vendor SAP announced that it was entering the SaaS market with a hosted version of its mySAP CRM product. The move allows a broader variety of businesses to take advantage of the company's offerings due to a more flexible, pay-as-you-go cost structure. Instead of the notoriously large up-front investments that SAP implementations are known for, the SaaS option now allows companies to pay for only the services they are consuming.

2.4. Multi-tenant Software Architecture

A building's structural architecture will differ depending on whether it is designed for a single occupant or multiple tenants. Similarly, software design will differ depending on whether it is intended to be used by a single "tenant" or multiple tenants. The term "multi-tenant" has therefore emerged in the context of SaaS systems to describe the design principles that allow the software model to be used by multiple unrelated parties.

One of the key distinctions between conventional software and SaaS software is that the latter has been designed since its inception to provide the data separation, concurrency, and manageability that are needed to deliver on the SaaS promise for multiple "tenants" using the same infrastructure at the same time. Thus, the software industry has borrowed the term "tenant" from the property management field to describe the architecture of the SaaS system, and why this architecture is necessary when looking at service providers.

The key architectural feature of multi-tenant software systems is the use of a single, common infrastructure and code base that is centrally managed and shared across all systems. This common infrastructure is what allows these systems to enjoy the economies of scale that reduce the total cost of ownership for all users of the system. Cost savings are achieved both through commonality of hardware resources, as well as sharing licensing expenses of operating systems and databases, to name a few, across a much larger population of users. (Again, this parallels the analogy to multi-tenant office buildings, where common expenses are spread across the entire group of tenants rather than simply absorbed by just one occupant.)

Multi-tenancy stands in sharp contrast to the practice of simply deploying old client-server architecture in a data center and calling it a “hosted service.” This approach to providing software services over the Internet traps service providers into a costly infrastructure where each customer is required to have their own server with the same cost structure as if it were still located at the customer site. Except for the fact that someone else is managing the system, this is hardly any different than the legacy approach to application ownership.

Even more problematic, however, is the attempt to serve multiple customers from within a single-tenant application that was not designed with sufficient data security to do so. A true SaaS application is designed to keep multiple customers each in “their own universe” so that they cannot see or manipulate each other’s data. This is clearly important for any type of application provided as a service, but particularly for the financial and security domains, not to mention the federal arena.

2.5. SaaS in Security Management

SaaS was introduced into the security industry in 2001 by Brivo Systems. The company provides a centrally hosted access control system that manages locally installed control panels, which are in turn connected to conventional card readers, door strikes, sensors, and the like. The system is also integrated with a variety of DVR/NVR systems to provide a single user interface for alarms, event history, and video that is accessible from anywhere on the Internet. Figure 5, below, illustrates this architecture:

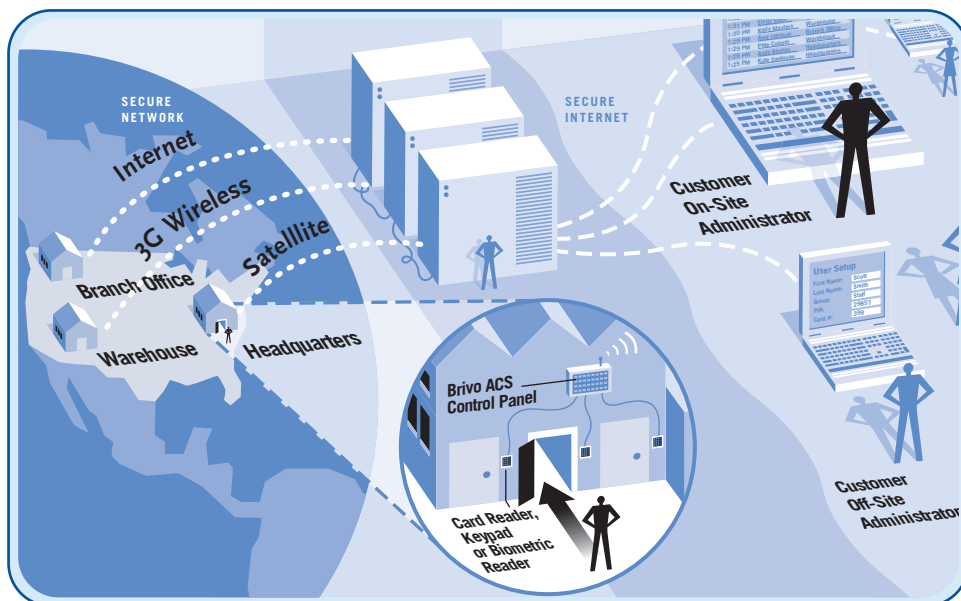


Figure 5: SaaS Access Control Overview

With regard to older PACS architectures, the primary innovation in this SaaS system is that it eliminates the requirement for a PC or server to be located at each property. Instead, it uses a centralized Web Service to perform all user and asset management, as well as all alerts, alarms, email notifications, and general reporting. Control panels communicate back to the central Web Service over a secure socket layer, or SSL, connection through any available IP communications channel, including LAN, broadband Internet, cellular, or even satellite.

Similarly, facility administrators gain access to their own accounts through the Internet. From their browser, administrators can perform all security management functions, including the option to:

- View event/alarm history
- View stored or live video
- Lock down facility
- Add/delete Credentials
- Add/delete/modify Users
- Add/delete/modify Administrators
- Add/delete/modify Doors/Readers/Elevators/Sensors

- Create/edit email notification rules
- Create/edit schedules
- View permanent journal of administrative actions
- Create and run reports

In short, the full set of system administration functions that are normally available only via the console of a client-server type of system are available to any authorized administrator via an SSL connection over the Internet.

2.6. Total Cost of Ownership

Total Cost of Ownership (TCO) is a well-studied discipline within IT at large, but its results have seldom been brought to bear on the world of electronic access control. Now that the Internet-based SaaS²⁹ model for access control management platforms is in many cases replacing the traditional server-based approach, property managers need to evaluate which is the most cost-effective solution for their organizations.

Until recently, it has been common for buyers to think of system costs in terms of one-time, up-front server and software expenses. However, recent studies have established that the largest part of application and server ownership costs actually exist in ongoing operational expenses, maintenance, and support agreements.^{30,31} This is particularly true of computer systems that provide infrastructure services such as access control, because they must be held to a higher standard of availability and performance than ordinary office equipment.

We find that for a typical “branch office” or “managed property” scenario, the SaaS model for security management platforms is the clear operational and financial winner, due primarily to the economies of scale introduced by hosted application services, as well as reduced up-front costs. The access control profile of such facilities includes:

- Base building access (lobbies, elevators, turnstiles, parking)
- Service doors, shared resources
- One or more tenant suites

One of the reasons that the SaaS solution is so much less expensive is evident in the contrast between the amount of installed equipment and servers required for a traditional approach, as shown in Figure 6, Legacy vs. SaaS PACS Infrastructure. Note that the server-based system architecture requires computing resources at each location, while the SaaS architecture centralizes this resource, yielding a scalable solution that can be shared across a large number of facilities.

Only 15% of the lifetime cost of server ownership is captured by the initial purchase price, which means that your \$1,000 server actually costs you over \$6,600.³²

²⁹ Note that during the history of centrally hosted software solutions provided as a service, SaaS has also been variously referred to as the ASP (Application Service Provider) model or, earlier, the MSP (Managed Service Provider) model. All three names refer to the practice of providing software to end users from a centrally hosted system, rather than from on-premises servers with installed software licenses. The significant difference between SaaS and both ASP/MSP models is that SaaS applications are generally required to be native, multi-tenant Web applications, rather than simply hosted client-server or other legacy applications with an HTML front end.

³⁰ “Total Cost of Ownership Reduction with VMware.” 2008, VMware.com http://www.vmware.com/vmwarestore/newstore/tco_login.jsp (March 10, 2008).

³¹ “SaaS—TCO,” op. cit.

³² “Total Cost of Ownership Reduction with VMware,” op. cit.

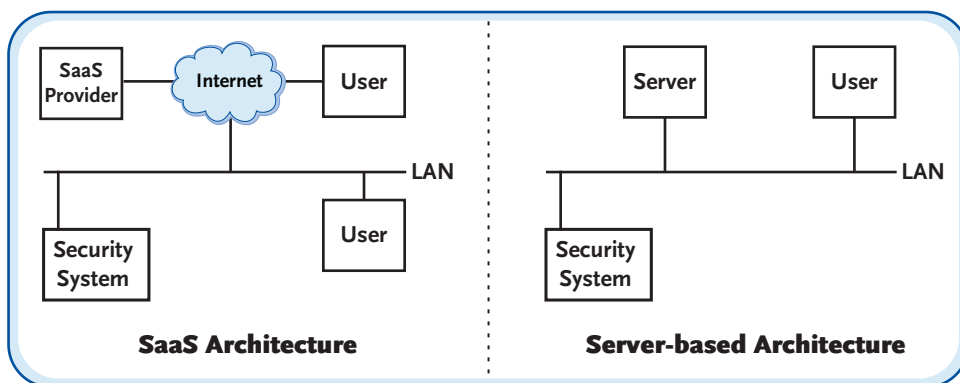


Figure 6: Legacy vs. SaaS PACS Infrastructure

The graph below shows the conclusions of a study from the point of view of a managed property, which demonstrated that for the systems the study modeled, the SaaS solution enjoyed an advantage of nearly \$26,000 (or 76%) over the server-based solution.

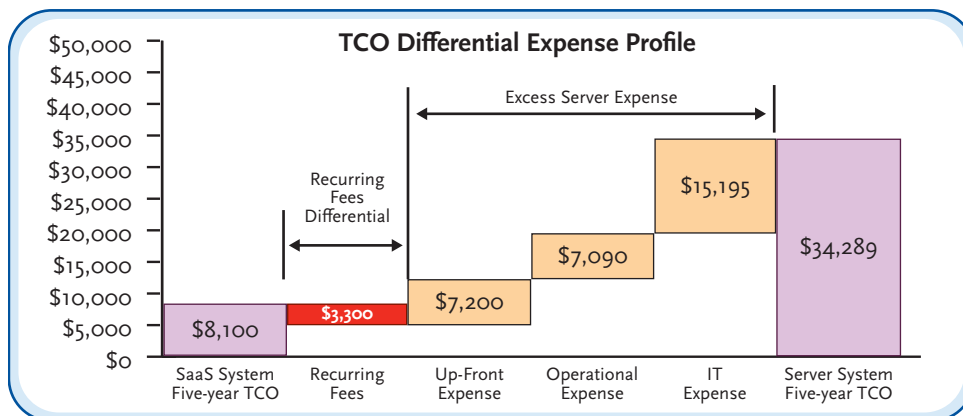


Figure 7: SaaS cost advantage over server-based system

The above chart shows realized savings without factoring in the IT expense or business impact of server downtime, which can be considerable for many types of enterprises. When these additional cost factors are added, the cost-effectiveness of SaaS solutions is even more dramatic.³³ As the full study on this subject demonstrates, these savings are multiplied with the number of facilities under management, with the same additional savings accruing for each additional property.

Tenant Suites

The savings for tenant suites are particularly large because the fixed cost of a control panel has already been absorbed by the base building access system, which means that tenant suites experience only the marginal cost of wiring a small number of additional doors. The information security provisions of the SaaS multi-tenant software architecture provide isolation of the tenant access permission structure, even while sharing physical and logical resources with the base building system.

³³ This section is an excerpt from the Brivo Publication, “SaaS—TCO: How Web-hosted Software-as-a-Service (SaaS) lowers the Total Cost of Ownership (TCO) for electronic access control systems.” Readers are referred to the full study for more details.

In addition, the fact that the base building access and tenant suites may choose to share a common credential database (e.g., FIPS 201 credentials) means that cardholder management can be performed in a single operation across both common and tenant-specific domains, thereby reducing lifecycle management costs provisioning. This avoids problems with multiple enrollments that are typically found when tenant suites use a different access control mechanism than the base building. It also reduces the number of credentials that employees and contractors must carry with them.

2.7. Robustness of SaaS

In addition to its other characteristics, SaaS solutions are often contrasted with distributed server solutions in term of their robustness—both day-to-day availability levels, as well as their ability to tolerate larger-scale disasters.

Availability

The availability of an online system is simply measured as the amount of time the system is performing correctly divided by the total time period being measured, usually expressed as a percentage. Because SaaS providers are contractually bound by Service Level Agreements (SLAs) to serve many customers at availability rates in excess of 99.9% or higher, SaaS providers have a very strong financial incentive to ensure that they meet these requirements. If they fail to deliver on this contract, many face stiff penalties under the terms of their SLAs.

The way that SaaS providers typically meet these high availability requirements is through system architecture and 24x7 technical staffing. System architecture approaches include:

- Redundant servers
- SAN storage arrays
- Hot-swappable RAID drives
- Remotely manageable servers
- Load balancing with dynamic rerouting of traffic
- Network monitoring

By way of contrast, most of the installed access control systems are single instances of PC or servers with none of the advanced techniques described above. Backup and redundancy are often either lacking, or must be provided by manual means such as driving to an affected property to switch to an alternate system.

That said, system architecture only goes so far toward ensuring high availability. Most SaaS providers also have trained, full-time staff dedicated to the maintenance and support of the services they operate. In addition to preventive maintenance to reduce the likelihood of problems, the staff is usually available 24x7 to respond to any early warning signs of potential issues. Most legacy security systems, by contrast, have no such dedicated staff to monitor their performance, and even under the best maintenance contracts, these older systems are usually 4 to 24 hours away from having a technician arrive on-site to investigate problems. This is not a standard of service that suits high-security applications.

Disaster Recovery

Disaster Recovery refers to the ability of an organization and its information systems to withstand and recover from a large-scale systemic disruption to its infrastructure. Typical scenarios used to judge the effectiveness of a Disaster Recovery strategy include studying the possible outcomes of natural disasters, such as hurricanes, or man-made disasters, such as terrorist attacks.

Because disasters on this scale may potentially affect a large area, such as an entire campus or city, the only effective preventive response to occurrences on this scale is to ensure that alternative information systems and infrastructure are in place at some large geographical distance from those they are backing up.

In the security industry, for example, a SaaS provider like Brivo Systems has multiple data centers throughout the United States to provide this level of insurance to its customers.

2.8. Information Security

A fundamental concern presented by many organizations considering SaaS is the security of the information contained in the provider's database. This is, of course, a prudent concern and one that every organization needs to consider when evaluating options for new information services.

That said, much has changed over the years since SaaS was introduced, and security has been demonstrated to be equal to or better than what most organizations are able to provide on their own. In the words of one industry analyst:

The introduction of software-as-a-service (SaaS) at the turn of the century was met with intrigue by the business value proposition and concern regarding the safeguarding of sensitive data over the wild, wild web by an outside organization. Several of the traditional on-premise software manufacturers did their best to fuel this fear and further surrounded security concerns with FUD (fear, uncertainty and doubt).

Seven years later the security concern has diminished for most—not because it is any less of a concern but instead because many SaaS vendors have demonstrated admirable security safeguards that go well beyond what most client organizations could achieve internally.³⁴

Today, SaaS organizations are usually flexible enough to provide a variety of security options around their services to meet any unusual needs that clients may have. This could include additional services such as restricted VPN offerings, separate servers, encryption of data-at-rest (in addition to standard SSL for data in motion), and biometric authentication, to name a few.

³⁴ "Fundamental SaaS Software Differences." 2007, CRMLandmark.com, <http://www.crmlandmark.com/saas-security.htm> (January 8, 2008).

3. SaaS and the Three Forces

Given that SaaS is now the dominant emerging computing model, and that it has been shown to provide significant financial and operational benefits for security management platforms, the question becomes: How does SaaS address the Three Forces shaping the federal security challenge?

In this section, we discuss specifically how the SaaS model for access control, as well as related services such as video surveillance, provide a model for complying with new mandates while saving both time and cost for property managers.

Our key conclusions are summarized below:

Force	SaaS Benefit
HSPD 12/FIPS 201	Facing widespread systemic changes to access control infrastructure due to HSPD 12 and FIPS 201, SaaS provides the most rapid and cost-efficient means of implementing compliant systems. It also provides higher degrees of uniformity across systems, and ongoing cost reductions due to administrative workflow reductions. Security is ultimately enhanced through the centralized management of employee data.
Physical/Logical Convergence	By reducing the number of additional identity stores required to manage federal employee and contractor access permissions, SaaS solutions simplify identity management burdens and hasten the process of physical/logical security convergence. Implementation costs are further reduced by cutting down on the number of system interfaces required.
FISMA	FISMA compliance for electronic security systems is simplified through centralization and by reducing the total number of systems to be accredited, monitored, and managed. The OMB has begun to recognize the broad impact of SaaS for improving FISMA compliance. Federal property managers can outsource significant compliance workload to SaaS providers and realize costs savings by leveraging a single compliance process across thousands of properties.

SaaS provides rapid, cost-effective adaptation mechanisms for all three of the forces shaping federal security management.

3.1. FIPS 201 and SaaS

It is a widely accepted conclusion that all of the PACS architectures currently providing electronic access control to federal facilities will need to be upgraded or replaced in order to comply with HSPD 12 and FIPS 201 requirements.³⁵ At a minimum, the PACS architectures must be able to accommodate the new card formats of the FIPS 201 PIV cards. Some legacy PACS architectures may be capable of doing this with firmware upgrades, but it is evident from industry discussions that many will not. Similarly, the software that provisions the control panels in a traditional PACS architecture (see [Figure 2, PACS above](#)) will also need to be upgraded or replaced, depending on the age of the system.

With this potentially wide-ranging overhaul of many facilities, the concept of implementing a unified solution based on SaaS technologies is very timely. The technologies that are chosen at this juncture will have to last for many years, through many changes to specifications and requirements, and at a low enough cost to be affordable under an “unfunded mandate.”

Implement One Change, Not Many

From an architectural, cost, and project management standpoint, the best system implementation is often the one that meets the requirements of the mission with the fewest components. As a centralized solution, SaaS provides agencies with the opportunity to put a single solution in place that covers millions of credential-holders and thousands of facilities. This greatly reduces system complexity and administration.

This simplified system design also translates into simpler in-building equipment installations, which means lower installation and maintenance costs. In a SaaS access control architecture, the only equipment installed on site is the control panel itself, eliminating one or more server or dedicated PC systems for each building under management.




Figure 6: Legacy vs. SaaS PACS Infrastructure

³⁵ For an excellent discussion of the full set of considerations for migration, see “Physical Access Control System Migration Options for Using FIPS 201-1 Compliant Credentials,” Smart Card Alliance Publication No. PAC-07002 (Smart Card Alliance Physical Access Council, 2007).

Scale of SaaS Systems Reduces Redundant Enrollment

SaaS systems in general—including those designed for security management and access control—are designed for massively larger user populations than their client-server counterparts. The reason for this design difference is present in the economic advantage that these systems create as a business proposition; that is, the scale advantage of providing a service to millions of users with a single application environment, rather than only a few thousand, as is normally the case. By the same token, a single SaaS system can also provide access control for tens of thousands of entry points—far more than traditional systems.

These scaling issues are relevant to the economic concerns of complying with HSPD 12 and FIPS 201 because of the sheer size of the federal employee and contractor populations, as well as the fact that these individuals must often access multiple properties. Clearly, this imposes not just a large administrative burden, but greatly magnifies risks of error and breached security.

Wide Geographic Reach

SaaS systems are also inherently designed to capitalize on the benefits and avoid the dangers of the Internet—including all of the security protocols needed to transact safely in that environment—a design consideration that allows SaaS systems to utilize the geographic reach of the Internet with no cost and no modifications to their core design.

For managers of federal facilities who are updating their access control systems in order to comply with HSPD 12 and FIPS 201, this geographic reach provides an opportunity to bring a dispersed set of buildings under a single security management policy that provides top-down oversight combined with the ability to delegate local control to local administrators.

One System, One Set of Updates

As with any system at the beginning of its technology lifecycle, there is a very high likelihood that standards, policies, and requirements for HSPD 12 and FIPS 201 systems will evolve over time. This has been true for every large-scale technology introduction.

In the absence of a system large enough to span very large user populations and very large sets of properties, an individual must be enrolled into the PACS of each and every building for which access is required.

3.2. Convergence and Identity Management with SaaS

Physical and logical security convergence is an emerging objective in many corporate settings. It brings about not only higher security, but also significantly reduced administrative burdens because a single workflow can be used to manage credentials for employees and contractors. Achieving this type of convergence, however, requires that an identity management solution be in place so that the multiple “identities” by which an individual may be known can be correlated with one another in order to control access across all of the relevant physical and logical resources.

There are now numerous identity management products on the market that can synthesize these multiple identities across multiple logical (IT) systems. These products usually also provide a Single Sign-On (SSO) capability that provides the necessary infrastructure for managing a user’s ability to gain access to any number of networks, software applications, etc.

The real issue for managing physical access control, then, becomes a question of the number of distinct access control systems and architectures that must be synchronized with the identity management solution—regardless of which particular technology is used for that solution.

For the federal agency working toward HSPD 12 and FIPS 201 compliance, physical and logical convergence presents a unique challenge because of the sheer number of systems across which identity—as well as other important access control information—must be synchronized. This situation is depicted in the diagram below with a traditional (one server per site) type of PACS architecture:

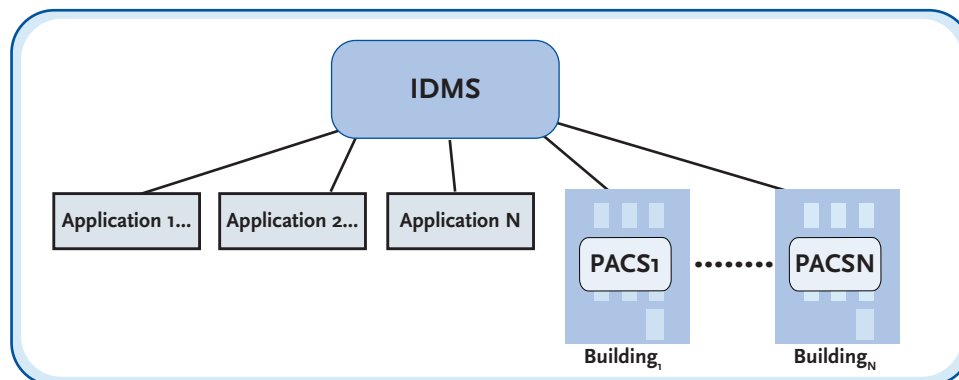


Figure 9: Identity Management Across Multiple Physical and Logical Domains

Here again, the SaaS architecture presents an opportunity for simplification of technical architecture and streamlining of administrative procedures, as shown in Figure 10, Identity Management in a SaaS Architecture. Part of this is due to a reduction in the sheer number of systems that must be managed; however, a larger part of the simplification is due to the fact that the data structures that define access control groups and their permissions—which are outside the purview of identity management

solutions—will be the same across an entire SaaS system, whereas older systems would require manual coordination across individually distributed systems and facilities.

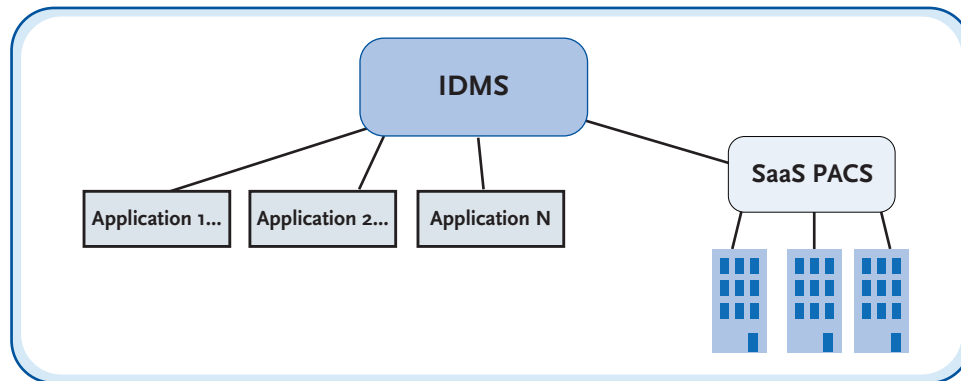


Figure 10: Identity Management with SaaS

The reduction of effort required to both implement and manage identities and their access control permissions for a single security system versus multiple systems cannot be underestimated. The next two subsections address the issues of administrative workflow and implementation difficulty in more detail.

Streamlined Workflow

To understand how administrative workflow burdens differ between a distributed, multi-server system and a centralized SaaS solution, consider what it takes to add an employee to a typical access control system. First, one must ensure that the value of the credential (card) has been entered into the system as a legitimate credential. Next, the administrator must enter a long list of personal information about the employee. Finally, the employee must be assigned to a group with appropriate privilege levels. This must be repeated for each facility to which the employee needs access.

Even under the best of circumstances where there is some degree of synchronization of user data across multiple access control domains, the entire process depends on having comparable or compatible mappings of access groups and the like across all of those independent systems. This degree of coordination and mapping across multiple systems installed at different times by different vendors under different sets of assumptions is very unlikely to be achieved, even when systems are first installed, let alone after a few years of operation when “drift” sets in and administrative policies at different locations begin to diverge.

By way of contrast, a SaaS solution offers the potential of centrally coordinating group access and permission structures according to an agreed template, and then enforcing consistency using that template across all subscribing facilities. With that degree of conformity already in place, the work of identity management, SSO, and other synchronization applications becomes much more powerful. Federal administrators can be assured that not only are identities being maintained consistently across all of these systems, but that the meaning of access control concepts like “contractor access privileges” have consistent definitions across all of these systems.

Reduced Interface Count for HSPD12 Ecosystem

One of the complications added to PACS architectures with the advent of identity management under FIPS 201 is that PACS will eventually need to communicate with a larger variety of other systems than in the past. The role of PACS architectures within the broader context of the “HSPD 12 ecosystem” has been outlined by the HSPD 12 Architecture Working Group.³⁶ A simplified version of their context diagram is provided below to illustrate how this affects the consideration of PACS implementation.

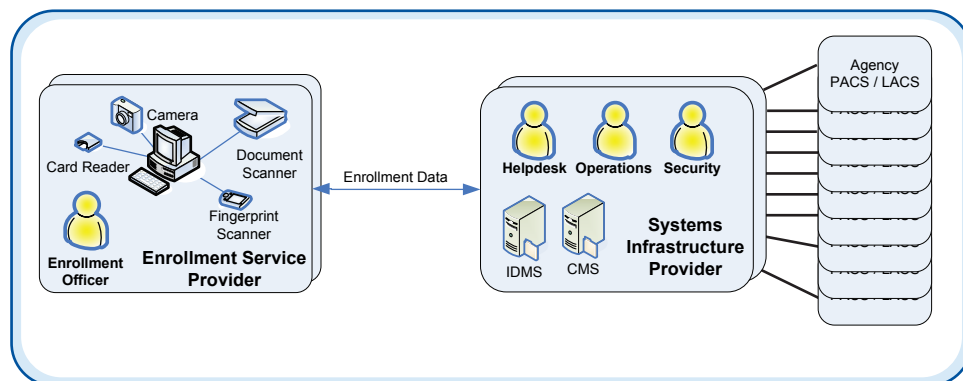


Figure 11: Traditional PACS Interaction with HSPD 12 Components

As shown in Figure 11, traditional PACS architectures will need to interface with other Systems Infrastructure Providers (SIPs) to receive data concerning credentialed users, identities, PIV cards, etc. In this traditional architecture, each individual PACS solution in each building would need to have an interface back to the SIPs—a complex and expensive proposition, considering that each one would have to be planned, designed, and validated for each installation.

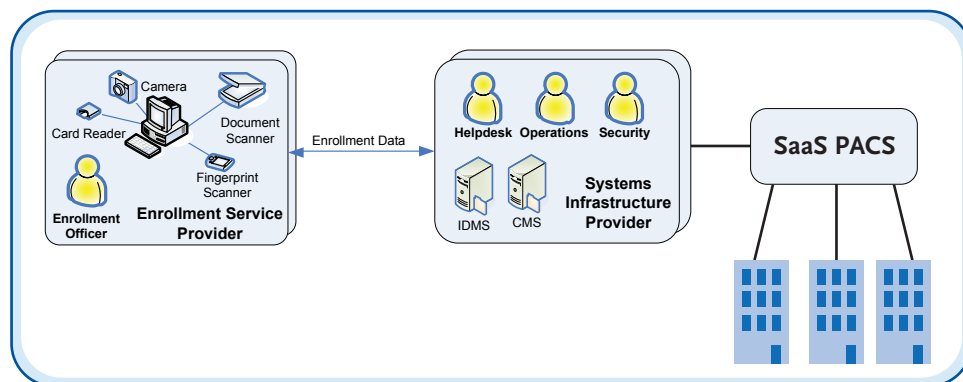


Figure 12: SaaS PACS Interaction with HSPD 12 Components

Contrast this with the SaaS option shown in Figure 12, where a single implementation between the SIPs and a SaaS provider can be shared across thousands of buildings. As the TCO study referenced above for general commercial properties, this drastically reduces up-front implementation time and expense. In this context, the SaaS solution again reduces complexity because the number of system interfaces is vastly reduced, as shown in Figure 12.

³⁶ The HSPD 12 AWG is convened under the auspices of the HSPD 12 Implementation Executive Steering Committee (ESC).

In the current environment where these interfaces between SIPs and PACS have not yet been fully defined,³⁷ it is of particular merit to keep their number to a minimum so as to limit changes when they are eventually defined.

3.3. FISMA with SaaS

Of the Three Forces shaping the future of access control and security management in the federal sector, FISMA is the one that has the broadest intersection with government IT initiatives in general. For that reason, it also has the most easily understood relationship with the movement toward SaaS as a preferred computing model for many applications.

Congressional Testimony on SaaS

In congressional testimony regarding information security and the future of FISMA, SaaS has emerged as a recommended solution to the many implementation difficulties presented by the new regulation. In particular, it is seen as a way for the government to save money while still maintaining the required information security standards. Speaking before committee in the Rayburn House Office Building during the summer of 2007, Paul Kurtz testified:

“Trends in the private sector, however, suggest a major paradigm shift. Many companies are migrating towards Internet accessible software and data services, which are often referred to as “software-as-a-service” (SaaS). Software-as-a-Service consists of applications and databases that are delivered to customers over the Internet from a shared IT infrastructure. By applying economies of scale to the development and operation of these applications, a SaaS provider can offer better, cheaper, more reliable applications than companies can provide themselves. Tens of thousands of businesses, including large financial enterprises, have already migrated sensitive data to SaaS providers and data warehouses, and it is projected that SaaS will account for nearly half of all software sales in the private sector within the next five years.

“Although federal agencies have been slow to adopt SaaS, I am confident that this migration to software and data services will eventually occur in the public sector as the government begins to recognize the cost savings and performance advantages of SaaS. This Committee should make sure that FISMA facilitates this evolution rather than hinder it.”³⁸

Support for SaaS within the OMB

The support for SaaS as a model for the government is not limited to private sector enthusiasts. Within the OMB, for example, Karen Evans, administrator of the Office of Electronic Government and Information Technology, is a strong supporter of the SaaS concept for government software initiatives.³⁹ She believes that the government needs to move to a more service-oriented model and shut down some of its legacy approaches to managing large projects. At a SaaS conference in Washington, DC, in January of 2008, she is quoted as saying, “We can’t continue to maintain all of the things we have. We have to start shutting down some of our legacy systems. We really

³⁷ Wynn, Bob. “Federal Security Mandates: An Update on HSPD-12 and FIPS 201.” May 15, 2007 Security Technology & Design.

³⁸ Kurtz, Paul B. 2007. “Federal IT Security: The Future of FISMA.” Paper presented before the subcommittee on Government Management, Organization and Procurement and the subcommittee on Information Policy, Census and National Archives of the House Committee on Oversight and Government Reform, June 7, in Room 2154, Rayburn House Office Building.

³⁹ Gross, Grant. “U.S. OMB Pushes for Software as a Service.” January 16, 2008, InforWorld.com, http://www.inforworld.com/article/08/01/16/US-OMB-pushes-for-software-as-a-service_1.html (January 29, 2008).

have to move to a ... service-oriented market.”⁴⁰ This perspective from the OMB was officially included in its 2007 FISMA guidance, serving as encouragement for agencies to use SaaS and other market solutions when it is possible to save money by doing so.

SaaS Addresses Many Requirements Off-the-Shelf

A quick look at some of the major categories of FISMA requirements versus SaaS characteristics helps to explain some of the recent enthusiasm for this model when it comes to government software procurement, as shown in the table below.

FISMA Requirement Category	SaaS Delivers
Media Protection	SaaS services routinely provide media protection in the form of multiple redundant databases, backups, and disaster recovery facilities. For reasons of client confidentiality, they also routinely sanitize or destroy worn or disposable media.
Physical Security	SaaS providers typically house their infrastructure in data centers with very high levels of physical security, including such measures as bomb/blast protection, photo and biometric ID, 24-hour guards, etc. Personnel access to information processing equipment is restricted based on predetermined access lists.
Environmental Protection	SaaS infrastructure is typically housed in environmentally controlled data centers that include redundant air conditioning, fire control, and backup electrical systems.
Incident Response	As a matter of doing business, SaaS providers have defined incident response procedures, along with staffing required to execute these procedures, thereby freeing the end user from having to maintain these resources and procedures.
Monitoring	Continuous monitoring against threats and failures is at the heart of every SaaS business, again providing one of the necessary elements of FISMA compliance on behalf of the end user.

Reducing Cost of Compliance

In each of the categories outlined above, there are structural reasons why SaaS providers are able to provide these services at a very competitive level compared to the expense of installing a slew of distributed systems to accomplish the same task. In most cases, the same class of service (e.g., environmental control or physical access restrictions) are already being delivered as part of the SaaS provider’s commercial offering. Doing so for a government client is, therefore, not an additional requirement, and does not need to incur any additional expense.

“With SaaS, you have someone who is an expert in that service and offers it to a wide variety of users.”

—John Murphy, USA.gov.⁴¹

⁴⁰ Gross, op. cit.

⁴¹ Miller, Jason. “Most Agencies Still Wary of SaaS.” February 4, 2008, Federal Computer Week, http://www.fcw.com/print/22_3/procurement/151479-1.html (February 10, 2008).

In this way, SaaS naturally lends itself to cost containment. The same economies of scale that apply to providing the services themselves also apply to ensuring the security of those services. This effect is recognized and in fact encouraged within FISMA regulation:

“Reuse and sharing of security control development, implementation, and assessment-related information can significantly reduce agency security costs in new acquisitions, certifications and accreditations of similar information systems, and reaccreditations of existing systems—and can ultimately result in a more consistent application of security solutions, agency-wide.”⁴²

Government Agencies Beginning to Use SaaS

While SaaS adoption is still early in the government at large, some agencies have begun to venture into SaaS for certain applications.⁴³

Agency	Application
DISA	Instant messaging, low-bandwidth text chat, Web conferencing and shared whiteboard services
EPA	Customer Relationship Management
Social Security	Customer Relationship Management
Census	Customer Relationship Management
GSA / USA.gov	Various Web analytics services
DoD	Collaboration tools through its Net-Centric Enterprise Services contract

The expectation is that agency SaaS usage will continue to spread as both the cost advantages and security are better understood by contracting agencies.

⁴² “Guide for the Security the Security Certification and Accreditation of Federal Information Systems,” op.

⁴³ Miller, op. cit.

4. Conclusions

Throughout this white paper we have examined how SaaS provides strong adaptive solutions to the Three Forces shaping federal security systems procurement. In summary, the primary advantages of using SaaS to meet the challenges posed by these three forces are:

1	SaaS systems provide a highly adaptable, centralized credential repository with low on-site infrastructure requirements to speed FIPS 201-compatible access control installations at a minimal cost per building and per user.
2	SaaS systems provide the open interfaces, ease-of-upgrade, and standards flexibility to adapt to changing identity management requirements, while reducing implementation expense by minimizing the number of interfaces and overall system complexity.
3	SaaS systems reduce the FISMA compliance hurdle by providing a centralized information architecture that keeps protected assets in a secure computing environment and minimizing the number of at-risk components that must be located on local area networks at remote facilities.

While these reasons are more than compelling enough, there are additional ancillary benefits to the SaaS model that go beyond the strict demands of the Three Forces that provide additional motivation to recommend this emerging model for cost-effective computing services.

4.1. Ancillary Benefits

In addition to providing a strong adaptive mechanism for the Three Forces, SaaS architectures for access control and security management also provide a number of important ancillary benefits, as described below.

Fewer Credential Databases to Maintain

The intrinsically multi-site span of SaaS-based physical access control systems makes it very easy for security managers to manage credentials, and places less burden and dependency on implementing a complete identity management solution.

Single Record of Personnel Movement

In commercial property management and other enterprises, a high value is placed on the strength of SaaS access control to provide a single record of personnel movement that spans all of the properties where the person has attempted to gain entry. Because all control panels in a multi-site system report back to a central database, SaaS provides a single portal where security managers can see all the activities of any person registered in the system. This is a capability that can never be duplicated by the legacy approach of having a separate system in each building.

SaaS Reduces Energy Consumption and Costs

SaaS also helps agencies to meet the requirements of Executive Order 13423 (January 24, 2007) which requires that agencies “improve energy efficiency and reduce greenhouse gas emissions of the agency, through reduction of energy intensity”. As noted earlier in our previously referenced study, SaaS systems exhibit demonstrable energy savings over equivalent distributed server systems, primarily because of more efficient utilization of computing resources. These savings are realized because a typical “enterprise class” server cluster for access control actually has the power to support a much larger building and user population. SaaS puts all of this capacity to work by deploying it across multiple facilities and user groups—savings both energy and cost in the process.

Lower Costs, Pay-As-You-Go

The SaaS “pay-as-you-go” subscription model eliminates the need for capital expenditures up front, dramatically reduces the risk of implementation failure, and creates a vested vendor relationship. With no need to install new hardware or software, implementations are accelerated, cheaper, and easier on the organization.

Anytime, Anywhere Management

Because SaaS applications are native to the Internet, they provide “anytime, anywhere” management of physical assets. Clearly, this is a convenience for busy property managers. More importantly, it can also provide better safety-of-life performance when changes must be implemented quickly or from remote locations.

Benefits for Tenants

When property managers implement SaaS security systems for base building access, they make it exceptionally easy for tenants to add compatible access control to their suites. Because these systems are browser-based, tenants do not need to add their own servers or dedicated PCs to get going. All they need is to add incremental door-specific hardware (readers, sensors, electric strikes). Often this can be achieved using expansion ports on the same control panel that is already in place for common area control. Best of all, because this sort of change does not affect information security, it does not trigger any additional FISMA review.

4.2. Alternative Ownership Models for SaaS

Stepping back for a moment, SaaS can be regarded as both an ownership or consumption model for software, as well as a set of architectural principles for building multi-tenant applications that can scale to very large numbers of users at a very low cost per user.

Throughout most of this white paper, the discussion of SaaS has focused on such offerings in the form that exists in the commercial world, where both of these aspects of SaaS are presented to the marketplace as one. However, it needn't always be that way—especially for sufficiently large property owners or managers, or for organizations (like perhaps parts of the federal government) that are subject to unique regulatory or compliance constraints.

For organizations that fall into these categories—especially large property managers—it may prove nearly as cost-effective to acquire and operate on one's own a SaaS-architected solution from a commercial SaaS supplier. This approach effectively gives an organization the same large-scale, multi-tenant applications needed to achieve economies of scale, while preserving such measures of ownership and control as may be needed to satisfy regulatory or other considerations. Of course, the buyer must also weigh the other advantages that come from a full SaaS offering: 24x7 staffing, automatic applications upgrades, continuous monitoring, and so forth.

A privately hosted (or dedicated) instance of a SaaS application is another alternative that falls somewhere between a fully outsourced SaaS solution and traditional licensing and ownership of software systems. Considered by some to be a 'best of both worlds' approach, this hybrid alternative allows certain customization and integration options that may not be possible with a fully shared system. This again is an alternative that most SaaS organizations would be willing to entertain for sufficiently large populations of users and facilities.

All of these choices present trade-offs that only the buyer can weigh, but they illustrate the flexibility available with the application architecture behind a well-designed SaaS system.