A Cybermodel for Privacy by Design

Building privacy protection into consumer electronics.

By Michael H. Davis, Ulrich Lang, and Sid Shetye

IMAGE LICENSED BY GRAPHIC STOCK

Does privacy protection matter in consumer electronics (CE)? What is privacy, how is it valued, and where does it sit in your organization today? Chances are, if you do not have a chief privacy officer or data-protection officer, your company is lacking in protecting critical data, let alone observing all the laws and statutory regulations dealing with privacy (e.g., audits, compliance, etc.). Managing privacy is crucial, especially considering the key mandated privacy requirements, such as those concerning personally identifiable information (PII), the Health Insurance Portability and Accountability Act (HIPAA), and the payment card industry (PCI). In addition, the privacy definitions and the policy and enforcement effectiveness are themselves varied and complex, and they change depending on where your data reside—i.e., the state, province, and country. For example, the European Union's (EU's) data-protection directive [1] is much stricter than the weak U.S. privacy laws. (Note that if you plan to market a global CE product, you should know about the Safe Harbor Framework.)

How does one start to protect critical data and observe the associated privacy requirements with many of the privacy rules and variables themselves in flux? Where common, ubiquitous privacy requirements are lacking, few (if any) implementation-level, definitive privacy specifications exist for developers to build privacy-enhancing technologies (PETs), including CE. Therefore, we collectively need a global privacy framework to design and measure capabilities; we chose the Privacy by Design (PbD) initiative [2] as an existing international effort to support. We developed a cybermodel that enables the PbD seven foundational principles (described in the "PbD Principles" section). The fair information practice principles (FIPPs) [3] are another set of high-level foundational requirements that are widely referenced and integrated in privacy rules and laws, as are the Organization for Economic Cooperation and Development (OECD) privacy principles [4], [5]. Both need to be accounted for in a cybersecurity for PbD (C4P) model. Thus, C4P

Digital Object Identifier 10.1109/MCE.2014.2361192
Date of publication: 17 December 2014

Privacy is a simple concept, but it is a complex endeavor to protect it.

will inherently address the key privacy-protection and control aspects from the start, making the actual data environment relatively agnostic to the ongoing global privacy environment churn.

Recently, the EU's top court's decision on privacy rights (against Google) added the right to be forgotten. We advocate that even the notion of PII (with 12 major attributes) or HIPAA's protected health information (with 18 key attributes) is likely not enough quantification for effective data characterization. There are hundreds to thousands of other identification attributes (from what you search or post) that can pinpoint you—is all the data and metadata being collected worthy of privacy protection? Privacy is a simple concept, but it is a complex endeavor to protect it. The protection must be provided from an enterprise view and start from the source, sensors, then devices and CE, and through the networks and ISPs, etc. This end-to-end (E2E) connectivity is where shared vulnerabilities are alive and prospering—especially from the lucrative cybercrime aspect. This article focuses on that enterprise privacy view, where anything that attaches to the net—especially CE devices—can have privacy built in—by design.

PRIVACY IN CONSUMER ELECTRONICS

CE devices store, process, and move sensitive data, a large part of which is either private user information or critical device-control parameters that need protection. The complexity of privacy protection is acerbated by the secure storage concerns with USBs, flash drives, and hard drives—along with the communication concerns with Bluetooth, mobile/cellular, text, etc. Thus, there is ample evidence that data protection and privacy must be addressed in CE, just as it needs to be for the multitude of Internet of Things (IoT) devices and other smart devices that improve the quality of life and productivity. We collectively need to address the data-security and privacy issues within a more global enterprise view, as this is the environment in which the IoT and CE capabilities operate. The principle of shared vulnerabilities means that a threat vector in one device or one location, through the massively connected environment in which we all operate, then affects us all as an added risk—especially to our privacy.

The essence of cybersecurity distills to trust and data protection. There are numerous ways both can be manipulated in an enterprise environment as well as in devices and CE products. To frame the CE cyber- and data-security concerns with privacy protection, several overall security issues with CE devices and the chips therein are provided as potential threat vector examples:

- side-channel attacks—including timing, protocol stuffing, power analysis, etc.
- register transfer language—security design and layout tools—potential security vulnerabilities
- embedded code/Android and Apple OS/firmware and HDL languages—inherent insecurity
- machine-to-machine communications, access controls/authentication (MQTT and/or Alljoyn)
- hardware-based security protocols (DRM, PEAR, EPC Geb2, etc.)—how well are they integrated?
- trusted platform module—the methods, specifications are there, but will people use them?
- chip-level encryption and key management—protection on the chip, including RF emanations
- minimize reverse engineering capability/forensics—protect your IP and minimize attack vectors
- physically unclonable functions (protecting your "crown jewels")—hardware primitives, etc.
- piracy/counterfeit chips/inserted malware—theft of your design, trade secrets, and methods
- complexity of design and integration (e.g., system on a chip)—communication pathways are geometric
- ghost circuitry, untrusted CAD tools, etc., where covert channels can be almost anywhere.

These concerns can largely be grouped under supply-chain risk management (SCRM), which every CE device producer should proactively practice in real time. That is besides the complex integration hurdles within multiple environments, with scalability and composeability challenges (within policy, protocols, interfaces, and standards).

PbD PRINCIPLES

Given the varying global privacy requirements, we developed our C4P model around the seven major principles in the existing international PbD initiative (whose creator is Dr. Ann Cavoukian), also mapping these seven principles to the 26 privacy controls in National Institute of Standards and Technology (NIST) 800-53a Appendix J [6]. Thus, our C4P model will inherently address the major privacy-protection and control aspects in PbD, FIPPs, and OECD, essentially encapsulating the key data-protection and control attributes and making them relatively agnostic to the ongoing vague privacy definitions and requirements churn. Current PET and CE products are generally device-centric and not integrated as part of an overall enterprise system of systems (SoS) architecture foundation. Hence, current privacy products and services cannot easily integrate into multiple environments or scale—in a continuum from one end device to another, likely different, end device.

The overall C4P technical approach and proposed specifications are described in the "C4P and OPF" section. The seven high-level PbD principles are listed here, with detailed cybercapabilities defined later to best engineer and operationalize a secure PbD:

1) proactive not reactive; preventative not remedial
2) privacy as the default setting
3) privacy embedded into design
4) full functionality—positive-sum, not zero-sum
5) E2E security—full-life-cycle protection

6) visibility and transparency—keep it open
7) respect for user privacy—keep it user-centric.

C4P MODEL DESIGN RATIONALE

The essence of our C4P approach is to develop an open privacy framework (OPF) using a services-based approach [similar to the platform as a service (PaaS) cloud construct] applying data-centric-security (DCS) methods, which are integrated into an SoS package using existing commercial off-the-shelf technology (COT). Our OPF foundation leverages, aligns, and is integrated with NIST's Risk Management Framework and Cybersecurity Framework. By developing and documenting a common OPF for which it is easy for PETs and CE devices to develop capabilities, C4P enables more integrated privacy capabilities to become available to enhance usability, reuse, and innovation.

The key elements of our C4P model are listed as follows, along with qualifying aspects and boundaries.

- Our hypothesis is that by endorsing and supporting PbD, we facilitate a structured discussion on effectively enhancing privacy protection overall, regardless of the environment.
- Data itself can be best protected by an approach covering the application layer to data store schema, where the DCS model most effectively describes that methodology. The related PaaS cloud model also works at the application layer, where C4P encapsulates the controls and data therein (similar to object-oriented programming methods).
- An implementation-centric C4P model would use privacy-community-acknowledged notional requirements set within a universal use case that is representative of the privacy ecosphere. We focus on the NIST 800-53 Revision A Appendix J as an initial privacy-requirements set. We also leverage the OASIS PbD Documentation for Software Engineers [7] and PbD-related privacy management reference model and methodology specification [8] (including their use-case template) to guide our C4P model.
- The C4P approach is built on top of a standard network environment with the typical information assurance (IA)/computer network defense (CND)/security suite. While the C4P protections use DCS methods, the infrastructure still needs its own security to ensure availability and overall system protection.
- The C4P principle added capabilities (+capabilities) shown in Figure 1 are as follows:
  - data security (DataSec) (with key management and access-control capabilities)
  - software/applications security (SW/AppsSec)
  - security policy architecture and design (PolicySec).

> The OPF EA can be used as a vehicle to help update the privacy specifications that developers need to build privacy capabilities into a collective cyber ecosphere, providing guidance for PETs and CE.

We suggest that any cybermodel supporting PbD must also integrate a security continuous monitoring (SCM)/security information and event management (SIEM) capability (SIEMSec) to monitor the infrastructure security posture and feedback to DataSec, SW/AppsSec, and PolicySec. For existing products that meet or exceed the C4P specifications defined herein, see the various company briefs at http://www.sciap.org/blog1/?page_id=1554.

The overall notional C4P approach is depicted in Figure 1, with capability details and technical specifications in the "C4P and OPF" section. A draft, high-level overall cybermodel for PbD with a much broader IT/risk view is available at http://www.sciap.org/blog1/wp-content/uploads/Privacy-by-Design-cyber-security.pdf.

IMPLEMENTING PbD

Any C4P approach to protecting enterprise privacy must do that in a common, fully integrated, easily executable, global manner. IT, cyber, and data must all be collectively designed, built, and operated from an E2E, SoS, fully harmonized approach. There is a natural hierarchy in our enterprise IT/network environment, where the major integration complexities generally arise in the

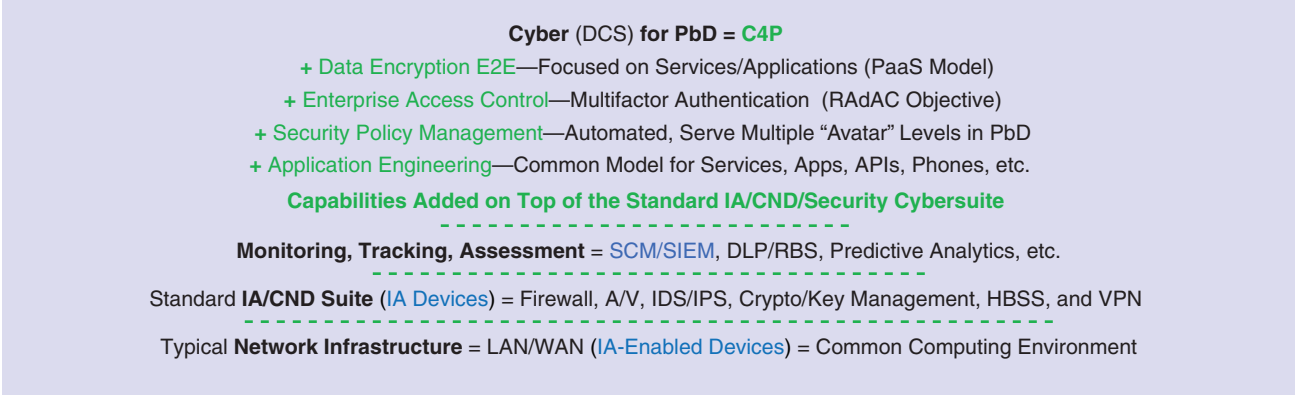

FIGURE 1. Building privacy protection into the enterprise—from the bottom up.

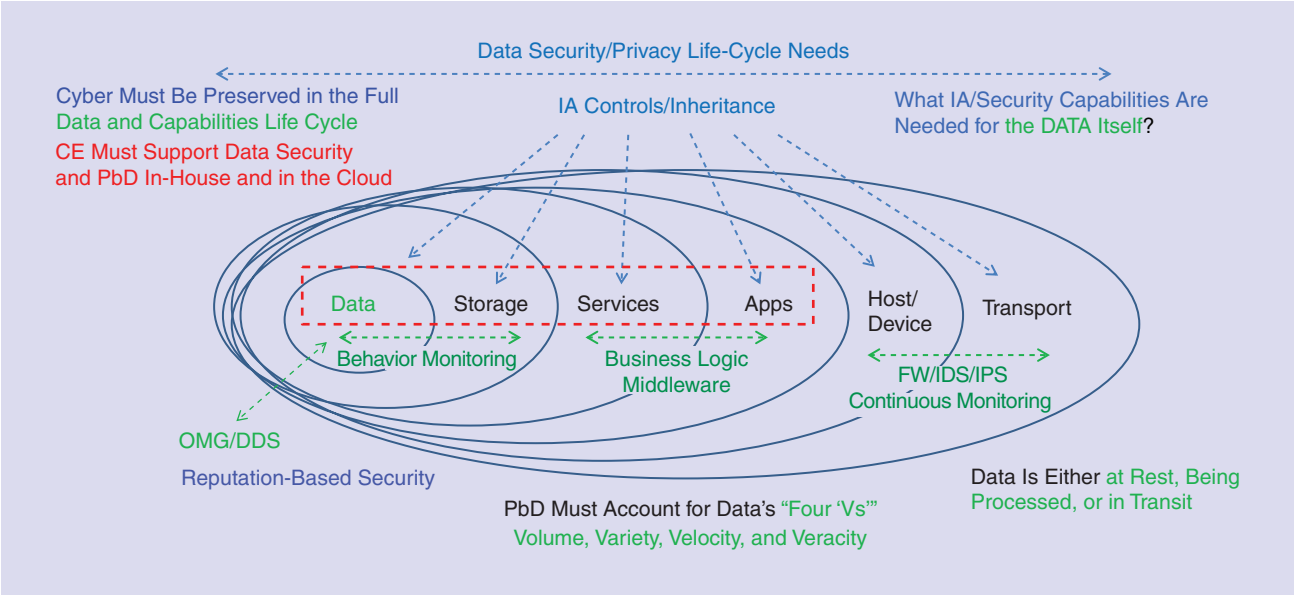

FIGURE 2. DCS—providing a defense in depth/breadth approach.

numerous interfaces, protocols, and many communications paths typically involved in E2E transactions within a SoS environment. The essence of privacy (data protection and access control) is in assuring the information-exchange requirements (IERs) between layers/enclaves and, specifically, the protections, controls, and inheritance aspects therein.

DATA-CENTRIC ARCHITECTURE AND SECURITY

Central to the overall execution efficiency and implementation consistency of our C4P approach is the general data-centric architecture (DCA) and supporting DCS approach to cyber. These methods are described in several illustrations and descriptions that follow, providing the technical perspective of our C4P model. Data-centric design recognizes that the essential privacy environment invariant is the protection and control of key data and IERs between systems or components. DCA describes the exchange in terms of a data model and data producers and consumers of the data. DCA/DCS relies on four basic principles: 1) expose the data and metadata, 2) hide the behavior, 3) delegate data handling to a data bus, and 4) explicitly define data-handling contracts.

In any C4P model, since privacy is all about control of assured data and IERs, we need to consistently account for how the data move throughout the enterprise and what the privacy protections and controls are at each layer. DCA decouples designs and simplifies communications and can

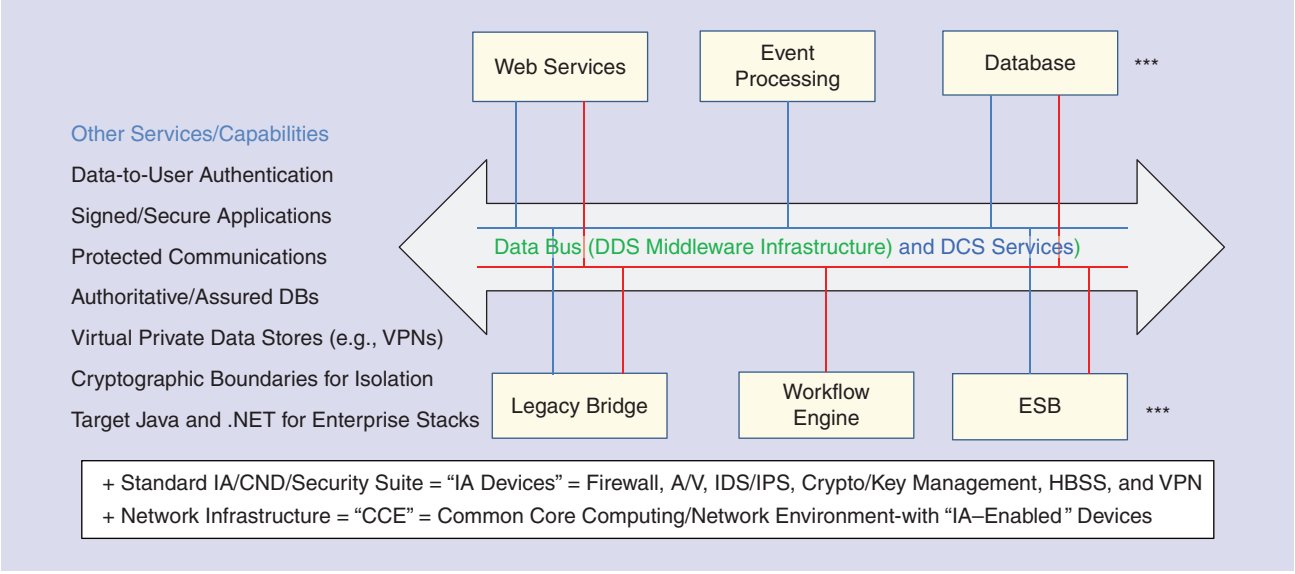

FIGURE 3. DCA and security overview.

link individual capabilities in an SoS environment into a coherent whole, using open standards—for example, Object Management Group Data-Distribution Service (OMG DDS) [10]—where details of the transports, operating systems, and other infrastructure information are then not essential to effectively implement DCS. Hence, C4P allows easier adaptation to performance, scalability, and fault-tolerance requirements. Figure 2 depicts a notional E2E DCA/DCS environment.

DCA AND DCS INTERRELATIONSHIPS

Within the DCA/DCS construct, we collectively need to quantify and modularize the key DCA components and capture the key security specifications (e.g., services, capabilities, and profiles). These include (but are not limited to): DCPS, DDSI, DataReader, DataWriter, Pub/Sub. Java, mobile code, widgets, storage functions, middleware, services, ESB, and, especially, application programming interfaces (APIs). API numbers will explode supporting IoT, with many billions of sensors and devices connected throughout the enterprise, including CE. The overall DCA/DCS interrelationships are depicted in Figure 3.

We propose that an optimum view to depict a C4P model is DCS added on top of the existing IT/network and IA/CND/security suite used in the typical enterprise environment (shown in Figure 3). As initially described, we view this C4P approach best as a services-centric, PaaS-like model, where the data, applications, and controls are encapsulated in an essentially "by-session or -transaction VPN" approach and thus agnostic to the overall infrastructure vulnerabilities. We account for the E2E access control and security policy details to satisfy the privacy elements in the "C4P and OPF" section. There are many benefits of a DCS approach within the PaaS model [11] that make it a very useful design methodology for any C4P approach.

Figure 4 depicts the various services-based models that we use to describe our C4P. The key point in the using the PaaS representation for C4P is that by protecting the applications and data layers, sensitive information is then inoculated from many, if not most, of the vulnerabilities in the lower layers. The overall systems availability still needs to be preserved by the infrastructure layers; thus, it is still essential that the typical IA/security/cybersuite be effective and maintained (and well monitored using SCM/SIEM).

> PbD must automatically provision the appropriate security controls and maintain appropriate threat and compliance monitoring as infrastructure environments scale up or down.

A C4P model has interoperability and composeability built in upfront (using DCA), as they help dramatically reduce complexity and ambiguity. Where used within a trusted cyberinfrastructure (TCI), this helps establish known risks, which reduces the attack surface, risks, and total operating cost; this TCI model with built-in trusted capabilities becomes the security infrastructure baseline for PbD. Our C4P approach also subscribes to the NIST "building in security" methods for a TCI, e.g., SP 800-160, *Systems Security Engineering: An*

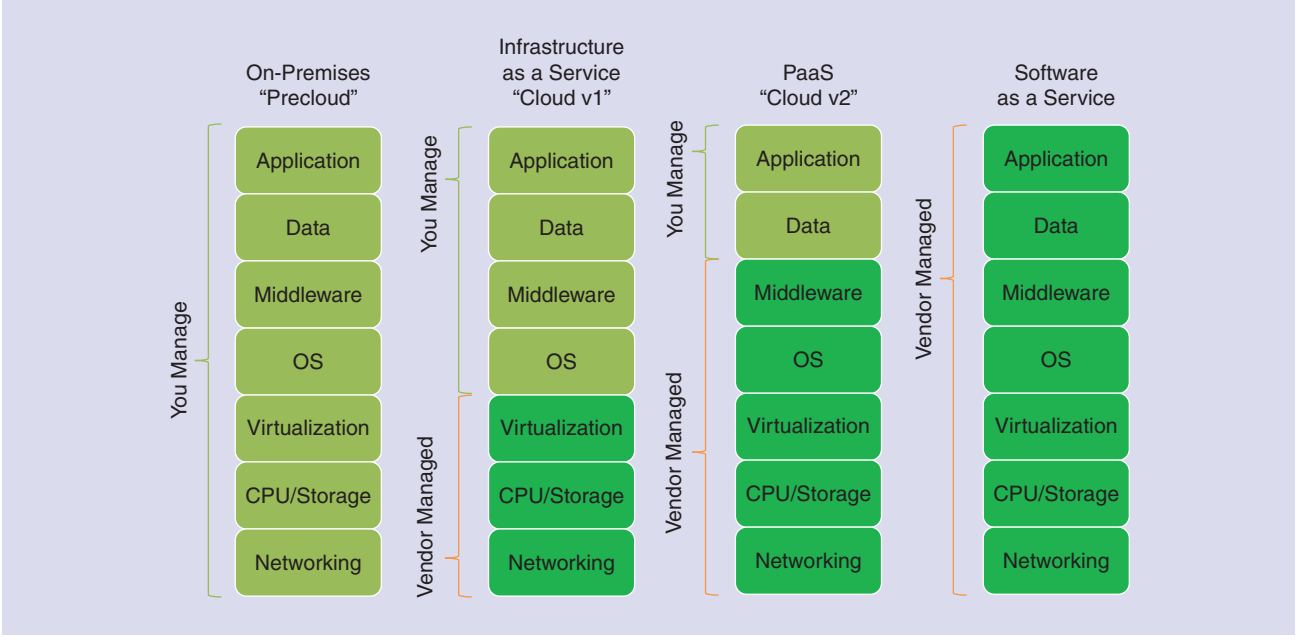

FIGURE 4. Various cloud views—highlighting the PaaS approach by C4P.

> The essence of our C4P approach is to develop an open privacy framework using a services-based approach applying data-centric-security methods, which are integrated into an SoS package using existing commercial off-the-shelf technology.

Integrated Approach to Building Trustworthy Resilient Systems (see http://csrc.nist.gov/publications/drafts/800-160/sp800_160_draft.pdf).

WHAT REALLY MATTERS IN CYBER

What really matters in implementing a common, affordable, and E2E TCI across all organizations? While the easy answer is it depends on the environment—there are several key cyber-factors we all collectively still need to build into our environments and better support PbD too. It is generally accepted that we have reasonably effective IA/cybertechnologies available now (to at least a first- and second-order effect). We all just need to integrate and maintain the cybersuite a lot better (principally using enforced cyberhygiene and effective access control).

We base our C4P model added capabilities and approach on the factors described so far—also subscribing to the four essential cybercapabilities that will endure, as called out in "Enterprise Software/SOA IA/Security Approach" [12]. The four main thrusts in the executive summary of the paper [12, p.2] are still germane now; thus, we embed these four technical aspects into our C4P approach:

- use IA/cyberstandards and related profiles within a functional cyberarchitecture
- provide E2E enterprise access control (integrity, authentication, and authorization) using an implementation centric approach (e.g., authorization-based access control)
- use a DCS approach (along with potentially adding content-based encryption)
- provide dynamic security policy execution among operations, management, etc.

C4P AND OPF

Each of the "+capabilities" functional details (listed in Figure 1: DataSec, SW/AppsSec, PolicySec, and SIEMSec) are described here in depth. This quantifies the support for the privacy capabilities stated in "Operationalizing PbD" (O-PbD) [13] specifically supporting the seven key principles listed earlier. A detailed explanation of C4P is necessary to provide a more complete description on how each capability fits each principle. We start the C4P story by using the O-PbD privacy definitions: "Privacy is about […] maintaining personal control over the collection, use, and disclosure of one's PII" (informational self-determination). Privacy is indeed a complex topic by itself, even without considering the differences between U.S. and EU policy. The O-PbD privacy perspective discusses many process related aspects of PbD as well: compliance, process improvement, and privacy policies need to be baked into applications across the entire software development life cycle. As stated earlier, the automation and embedding of security policy to manage the controls needed in various privacy levels and environments is the most critical aspect of any cybermodel for PbD. We all must, of course, build any cybermodel to the same privacy requirements, ideally quantified in privacy specifications derived from common, approved, authoritative privacy sources.

DATASEC

DataSec (E2E encryption, data-centric services, key management, and access control) capabilities for PbD include the following.

- *User security*: PbD requires that only authenticated and authorized users have access to the privileged parts of their PbD enabled applications. To restrict access to other users, DataSec provides multifactor authentication, which covers location, time, biometrics and other sensor data from the user before allowing access to the more sensitive parts of a PbD-enabled application.
- *Security against data breaches*: Data breaches are now a routine occurrence, and with the proliferation of cloud computing, much of the sensitive data processed by PbD-aware and -enabled applications is now resident in cloud data centers. A PbD application must have data-in-transit and data-at-rest security at the back end but also allow the trust footprint to be smaller. This means that the database servers, file servers, administrators, data-center technicians, or any intermediate equipment can all be untrusted.
- *Better operational awareness*: Security exceptions for both user security and data security are logged for audits, and outlier events raise alerts to users and application owners.

SW/APPSEC

SW/AppSec (Apps/Services and Phone/Mobile) capabilities for PbD include the following.

- *Automated*: PbD dictates for the need of automated policy authoring, enforcement, and auditing. If the security is based on manual processes, then points of error, vulnerabilities, and noncompliance are likely to be created.
- *Ubiquitous*: Omnipresence—the same control and management implementation should be operable on any environment, regardless of physical location, operating system, virtualization platform, or deployment method used. Policies defined by PbD should accommodate all the entities (hardware and software) and their operation (message exchange, file storage, etc.) within the environment.

- *Scalable*: The system should automatically grow and contract to meet the changing demands of applications and underlying infrastructure. PbD must automatically provision the appropriate security controls and maintain appropriate threat and compliance monitoring as infrastructure environments scale up or down.
- *Multilayer visibility*: Privacy challenges exist in both hardware and software. Any SW/AppSec solution considers privacy as an integral part of security and, hence, provides comprehensive solutions at each operational hardware and software layer.

> C4P will inherently address the key privacy-protection and control aspects from the start, making the actual data environment relatively agnostic to the ongoing global privacy environment churn.

POLICYSEC AND SIEMSEC

PolicySec and SIEMSec (security policy architecture and SCM/SIEM) capabilities for PbD are as follows.

- *Policy authoring*: PbD needs an intuitive, user-centric privacy policy authoring feature for users to set their privacy policies (informational self-determination). PolicySec must provide more functions than just enforcement of access-control policies.
- *Policy enforcement*: PbD needs a tool that maps these intuitive privacy policies into technical enforcement (access control, confidentiality, etc.) across the information life cycle and software development life cycle as well as configurable privacy code libraries. Model-driven security (MDS) is a good tool for such a mapping [14]. Attribute-based access control (ABAC) and encryption are example mechanisms that can be configured to enforce the privacy policies.
- *Policy audit*: PbD needs a user-centric tool that lets users verify (audit) that their policies are enforced correctly. PolicySec and SIEMSec help audit as-is processes and controls against the defined security policies for privacy.

MDS policy automation is the tool-supported process of modeling security requirements at a high level of abstraction and using other information sources available about the system (produced by other stakeholders). These inputs, which are expressed in domain-specific languages (DSLs), are then transformed into enforceable security rules with as little human intervention as

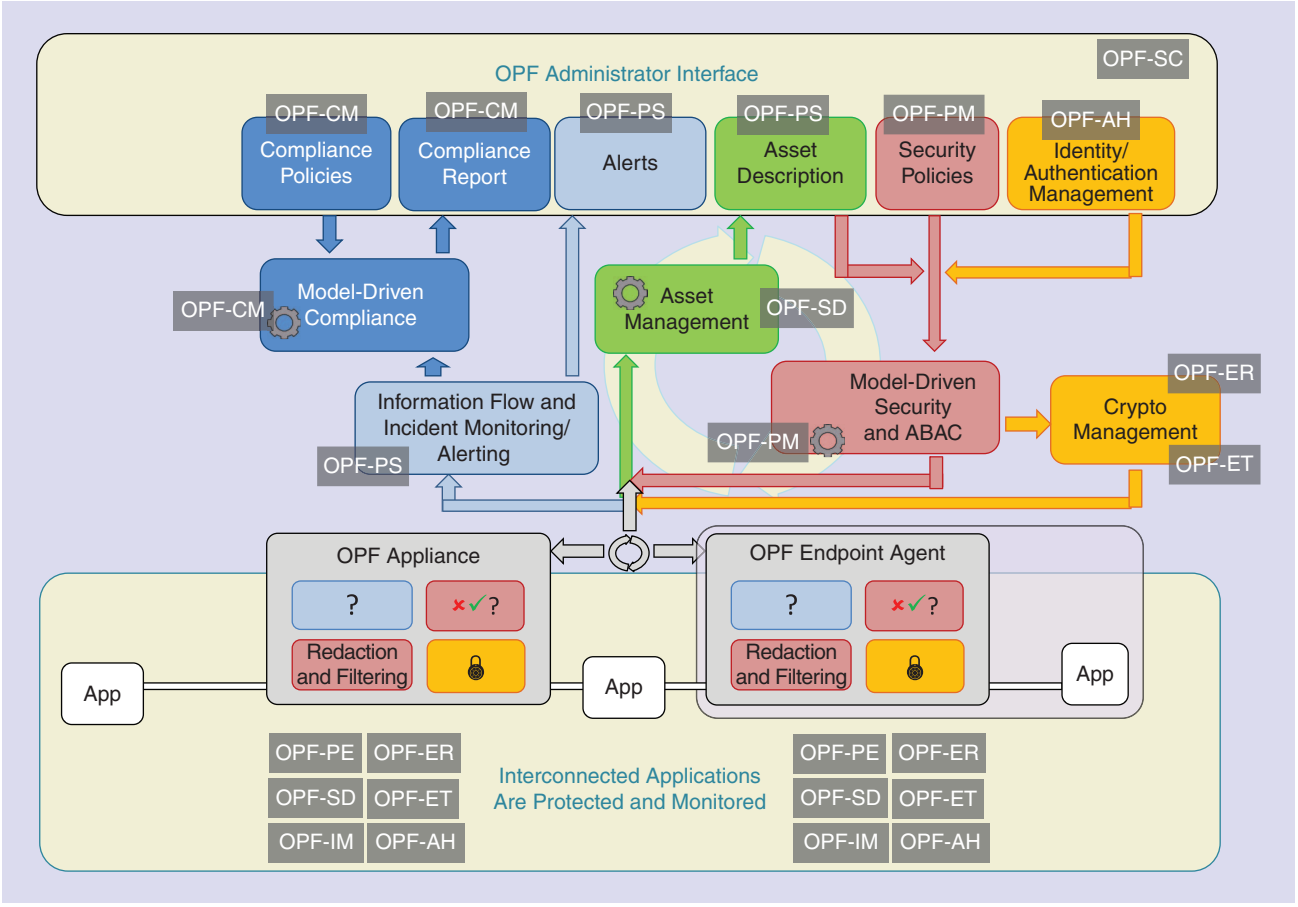

FIGURE 5. The reference implementation combined architecture [15].

The essence of cybersecurity distills to trust and data protection.

possible. MDS explicitly also includes the run-time security management (e.g., entitlements/authorizations), i.e., run-time enforcement of the policy on the protected IT systems, dynamic policy updates, and the monitoring of policy violations.

Model-driven security accreditation automation (MDSA) automates the analysis of traceable correspondence between technical security policy implementation (e.g., ABAC) and the IA requirements captured in undistorted requirements models (e.g., common criteria, control objectives). MDSA also documents supporting evidence for accreditation based on various information (esp. design-time system/security models, system/security artifacts, system/security model transformations, and runtime system/security incident logs). Furthermore, MDSA enables the automated change detection and analysis to determine whether the accreditation is still valid.

INFORMATION LIFE-CYCLE VIEW

The essence of our C4P approach is to develop an OPF using a service-based, PaaS-like approach applying DCS methods, which are integrated into an SoS EA foundation using existing COTS products. The C4P high-level reference implementation management approach, shown in Figure 5, includes functional design specifications for each OPF-xx capability shown. Each function includes a traceability aspect to assure compliance with all requirements through the product life cycle—as well as detailed specifications for each that allow developers more definitive build guidance. The diagram and related descriptions are based on research projects described at http://www.ict-icsi.eu/description.html and http://www.valcri.org/.

Privacy protection can be viewed as an information life-cycle governance/management problem. Privacy policies need to be enforced for information, including collection/creation, access (including delegation), transmission, storage, redaction, deletion, expiration, etc. In addition to numerous nontechnical controls (discussed in the O-PbD literature), a number of technical information security features need to be implemented. Our reference implementation technical approach is built on the unique combination of a set of common technical components, for which reference implementations are, for the most part, already developed and deployed by the partners. The sum of our integrated OPF enterprise architecture enabling PbD is much greater than its component parts as it also reduces the fog of privacy requirements (e.g., being infrastructure agnostic). This simplifies the overall privacy ecosphere and facilitates software and applications developers, PETs, and CE companies in developing interoperable privacy capabilities.

The main technical components and functions in the information life-cycle reference implementation are as follows.

- *OPF-PM: Policy Management*—PbD needs a manageable, intuitive, user-centric privacy policy authoring feature for users to set their privacy policies (informational self-determination) governing users, systems, applications, and interactions (information flows). It needs to allow users and administrators to author and/or select privacy policies captured in intuitive models (OMG-style domain-specific languages, DSLs).
- *OPF-PE: Automated Security Policy Enforcement and Alerting*—PbD needs a tool that enforces technical privacy rules and configurations generated by OPF-PM technically (access control, confidentiality, etc.) across the IT landscape (multiple layers of the system/application/network/VM etc.), across the information life cycle and software-development life cycle.
- *OPF-CM: Compliance Management and Automation*—PbD needs a user-centric tool that lets users verify (audit) that their policies are enforced correctly. This feature analyzes the traceable correspondence between technical security policy implementation (e.g., ABAC) and the IA requirements captured in undistorted requirements models (e.g., common criteria, control objectives).
- *OPF-SD: SoS Discovery*—The system automatically generates a model of the enterprise networks, systems, applications, information flows, users, etc. This system description plays a similar role as common criteria's target of evaluation.
- *OPF-IM—Incident Monitoring:* The solution needs to be able to watch network activity (including bandwidth usage), access control incidents, and more, by automatically capturing and analyzing anomalies detected in PbD appliances and/or locally installed policy enforcement point (PEP) software proxies.
- *OPF-PS: Presentation of (Current) Status*—The solution displays the current privacy posture on a continuous basis in a consolidated fashion. This includes the network status (e.g., in a Web browser or its 3-D asset viewer), a dashboard that reports on levels of events with options to drill into details—even the triggering network packet, a policy incident viewer, a compliance evidence viewer, etc. These events need to be categorized and graphed to display the state of the SoS.
- *OPF-SC: Security Administrator Collaboration*—The solution also includes a way for administrators to collaborate to resolve issues (e.g., a secure social network to facilitate collaboration between administrators).
- *OPF-ER: Encryption for Data at Rest*—All cryptography is configured and managed in a unified way together with the other policies in OPF-PM. The cryptography should be at least U.S. National Security Agency (NSA) Suite B certified and should not only encrypt the data for privacy but should also have checks for data integrity.
- *OPF-ET: Encryption for Data in Transit*—All encryption is configured and managed in a unified way together with the other policies in OPF-PM. Data in transit between storage and processing or between processing elements may be protected by SSL for an outer layer of encryption but must have an inner layer of encryption to be protected, similar to

the provisions in OPF-ER. It must therefore preserve the OPF-ER requirements in terms of NSA Suite B cryptography, privacy as well as integrity checks, security partitions, access-control lists, audits, and prioritized alerts.

- *OPF-AH: User/Machine Authentication*—User authentication can be based on multiple factors, e.g., the user password or PIN, a cryptographically secure time-based one-time password or token, successfully matched facial patterns of the user, the location of user, and the time of the request by the user.

SUMMARY

Privacy matters everywhere, the IoT and CE included, and the downside of not protecting privacy is greatly increased individual, company, and organizational risks, which are costly and last forever in the public domain. While we focused on the enterprise privacy-protection methods herein with our C4P approach, the basic elements of trust and data protection must be applied at the end devices as well, supporting E2E privacy. The threat vectors are too numerous to chase and try to fix/patch individually (and they morph and change by the hour); thus, privacy must be built in by design from the start, from the end points through the ISP. It does little good to fix the SCRM aspects mentioned at the beginning for CE devices, when what you connect to is vulnerable—as privacy is still lost. In the legal world in which we live, cyber third-party suits will become the norm, especially with data breaches; thus, a poorly secured CE device is a huge financial risk. Just as you must follow the UL rules, so will you need to do for privacy.

Our proposed C4P model and accompanying OPF EA is an executable foundation on which to build interoperable secure privacy capabilities into any standard IT network environment with the typical IA/cybersuite. This enterprise view starts from the sensor data methods, through devices/CE, the network protections, and connections to the Internet. The OPF EA can be used as a vehicle to help update the privacy specifications that developers need to build privacy capabilities into a collective cyber ecosphere, providing guidance for PETs and CE. C4P provides an E2E, SoS, and OPF that can scale, adapt, and endure and work well in most environments. In addition, this C4P model can be implemented now, while the global privacy requirements and the rest of the technical world catches up to making privacy a priority, thus minimizing long-term privacy liabilities and costs in the interim for users and companies alike.

ABOUT THE AUTHORS

Michael H. Davis (mike.davis.sd@gmail.com) received his electrical engineer and M.S.E.E. degrees from the Naval Postgraduate School, and holds CISSP and CISO certifications. He is a cybersecurity and risk management consultant. He is a cofounder and CEO of ACME Cyber Solutions. He has more than 25 years of experience in IT/IA/Cyber technical and operational leadership positions in many diverse government and commercial programs/venues throughout the product/service life cycle.

Ulrich Lang received his Ph.D. degree from the University of Cambridge Computer Laboratory (Security Group) on conceptual aspects of middleware security. He is a cofounder and CEO of ObjectSecurity. He is a renowned thought leader in model-driven security, access control policy, Cloud/SOA/middleware security, and the Internet of Things. He is on the Board of Directors of the Cloud Security Alliance (Silicon Valley Chapter) and works as a technical expert witness.

Sid Shetye received his master's degree in electrical engineering from the University of Southern California and his M.B.A. degree from UCSD's Rady School of Management. He is the CEO and founder of Crypteron. He is a leading security expert with more than ten years of experience in security and cloud software. He worked in management and engineering roles at Qualcomm, Broadcom, and RoboCFO.

REFERENCES

[1] Data Protection Directive 95/46/EC. [Online]. Available: http://en.wikipedia.org/wiki/Data_Protection_Directive

[2] A. Cavoukian. PbD: Privacy by design—Seven principles. [Online]. Available: http://www.privacybydesign.ca/index.php/about-pbd/7-foundational-principles/

[3] Fair information practice principles (FIPPs). [Online]. Available: http://www.nist.gov/nstic/NSTIC-FIPPs.pdf

[4] Organization for Economic Co-operation and Development (OECD). Guidelines on the protection of privacy and transborder flows of personal data. [Online]. Available: http://www.oecd.org/sti/ieconomy/privacy.htm

[5] (2013). OECD privacy principles. [Online]. Available: http://www.oecd.org/internet/ieconomy/oecdguidelinesontheprotectionofprivacyandtransborderflowsofpersonaldata.htm

[6] P. Gallagher. (2013, Apr.). National Institute of Standards and Technology. NIST 800-53Rev4—Appendix "J"—Privacy controls. [Online]. Available: http://dx.doi.org/10.6028/NIST.SP.800-53r4

[7] A. Cavoukian and D. Jutla. (2014, June). OASIS privacy by design documentation for software engineers (PbD for SW SE) (DRAFT!). [Online]. Available: http://docs.oasis-open.org/pbd-se/pbd-se/v1.0/csd01/pbd-se-v1.0-csd01.doc

[8] J. Sabo and M. Willett. (2013, July). OASIS—Privacy management reference model and methodology (PMRM). [Online]. Available: http://docs.oasis-open.org/pmrm/PMRM/v1.0/cs01/PMRM-v1.0-cs01.doc

[9] National Information Assurance Partnership (NIAP)—Product compliant list. [Online]. Available: https://www.niap-ccevs.org/CCEVS_Products/pcl.cfm and https://aplits.disa.mil/processAPList.action

[10] Object Management Group—Data-Distribution Service (OMG DDS). [Online]. Available: http://portals.omg.org/dds/

[11] D. Orlando. (2011, Jan.). Platform as a service (PaaS) cloud model. [Online]. Available: http://www.ibm.com/developerworks/cloud/library/cl-cloudservices2paas/

[12] A. Budgor and M. Davis. (2008, Oct.). Enterprise SW/SOA IA/security approach (clarifying the fog of IA). [Online]. Available: http://www.sciap.org/blog1/wp-content/uploads/Enterprise_Software_IA_Security_approach.pdf

[13] A. Cavoukian. (2012, Dec.). Operationalizing PbD (O-PbD). [Online]. Available: http://www.privacybydesign.ca/content/uploads/2013/01/operationalizing-pbd-guide.pdf

[14] Model driven security. [Online]. Available: http://en.wikipedia.org/wiki/Model-driven_security

[15] U. Lang. (2012, Nov.). Press release—ObjectSecurity is awarded multi-year "ICSI" intelligent transport systems EC FP7 research contract. [Online]. Available: http://www.objectsecurity.com/doc/20121101-icsi.pdf