



THE SECURITY STACK

A MODEL FOR UNDERSTANDING
THE CYBERSECURITY WE NEED

A WHITE PAPER



ABSTRACT

This paper proposes a four-layer model called the “security stack” as a means to visualize the complexity of cybersecurity problems and see through to comprehensive, effective solutions. The authors use the term “stack” strictly as analogous (having similarities) to other well accepted stacks (e.g., the OSI model) where layers deliver services and exchange information to achieve a higher level service. The notion of a “security stack” serves the proposition that security must be an integrated set of services. The paper defines each layer, offers examples of enabling technologies, related standards, and types of professional security services that implement the enabling technologies. It also notes where adequate enabling technologies or standards still need to be developed, or where policies need to be set and implemented to allow information to be exchanged between layers fast enough to keep up with the speed of emerging threats.

The security stack, furthermore, is consistent with the idea that information-communications technologies (ICT) must be architected and that security is a vital element in the ecosystem of ICT architectures. Just as multiple blueprints (electrical, plumbing, flooring, etc.) are required to construct a safe and stable building, the cybersecurity blueprint is an integral part of an ecosystem in which ICT architectures are made secure and sustainable by design – that is, intrinsically secure. In closing, the authors describe a variety of benefits of the security stack, which include serving as a guide for integration efforts, creating forcing-pressures for collaboration and, most importantly, establishing that no single layer can effectively contend with the sophisticated attacks of our present day and into the future.

INTRODUCTION – WHAT IS A SECURITY STACK AND WHY DO WE NEED ONE?

Efforts so far to address the challenges of cybersecurity have involved haphazard approaches with add-on technologies integrated through the best efforts of end users. These approaches no longer work, if they ever did.

Modern societies have become overly dependent on cyber systems¹ without adequate protections. Many of the tools and supply chains constructed to support human activity globally are interconnected with and interdependent upon cyber systems that are insufficiently secure. There is an imbalance between benefit (e.g., ease of use) and risk that is international in scope and byzantine in complexity.

It matters greatly that we get to a better equilibrium of benefit and risk, understanding the costs and implications. The risk associated with cyber systems goes straight to the heart of physical systems for transportation, energy, finance and health, to name just a few critical areas. We can and should solve these problems – now. We can and should put in place a framework to identify these problems more effectively today and manage them for tomorrow. That is the justification for introducing the concept of a “security stack.”

The notion of a security stack draws from other models defined as stacks not in the purest sense, but rather analogously. For example, the Internet Stack or the Open System Interconnection (OSI) models served to describe layers of services that interconnect by passing and receiving information to and from adjacent layers. Taken together these layers help define a functional computer network delivering some application or higher-level service. Models serve the same purpose as diagrams on the proverbial paper napkin. They are aids to simplify; to divide and conquer; to help in understanding large, complex problems by decomposing them into smaller, discrete components (layers in this case), which is the purpose of the security stack.

In many respects, the security stack is also a framework for the future. We are not lacking for a wide range of component security technologies, but they generally are used independently. The security stack helps visualize the value that can be derived from integrating these separate technologies in the same manner the OSI model serves as visualization for future network integration.

Efforts so far to address the challenges of cybersecurity have involved haphazard approaches with add-on technologies integrated through the best efforts of end users. These approaches no longer work, if they ever did. Despite the billions of dollars of commercial and government funds spent on them annually, the evidence is clear that risks grow and their impacts magnify. Add-on security devices that overlay network operations to protect complex systems are insufficient for the task, because they do not address vulnerabilities that reside within the operational network. This is why we have “patch Tuesday” to mitigate for vulnerable code, why worms like Sasser took root in 2004 exploiting buffer overflows, why Conficker

¹ A cyber system is defined to encompass all ICT systems (including hardware, software) operating and dependent on network connectivity.

can take command of vulnerable computers, and why advanced persistent threats constitute serious dangers to corporations and governments.

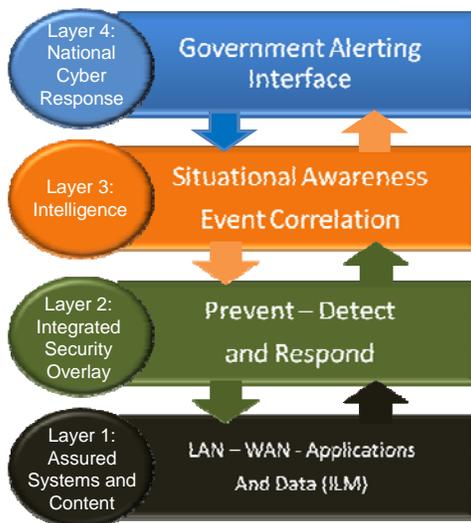


Figure 1: The Security Stack

Even the Internet has human history to contend with, a history where it has been necessary to protect land space, sea-space, air-space, and space-space (if you will). Now societies must consider the defense of cyberspace.

Security overlays have their value to be sure, but they are not enough. Systems need to be made secure from the start of their design; hence the term “security-by-design” was coined. Consider then the first two layers of the security stack, beginning with Layer 1, *Assured Systems and Content*, and followed by Layer 2, the *Security Overlay*. Both these two are necessary but together still not sufficient. We need more protection to complete the security task.

We also need situational awareness. We need to see the world of cyberspace both inside and outside our computing enclaves. This happens in Layer 3 – the *Intelligence Layer* – which correlates information from sensors to give advance warning of threats. The Intelligence Layer detects threats so that defenses can be adjusted, ports closed, and mitigations enacted before attacks can achieve their intended purposes.

We also need a Layer 4. Real-world borders of national sovereignty must be protected through cybersecurity. In this statement resides a paradox: We must defend physical, politically structured borders even though cyberspace has no borders, at least not in the sense of the geographic boundaries of nations. Cyberspace cannot divorce itself from physical space, as we are physical beings. The Internet’s pioneers envisioned a cyberspace to be free of old conventions, but such idealism, though it persists, is at odds with reality. We need the means in cyberspace to counter an adversary intent on doing harm to interests defined by the geographical boundaries of nations. Those means reside in our fourth layer – call it the *National Cyber Response Layer* – which represents the ability to protect the sovereignty of countries whose governments are pre-eminently concerned with the safety of their citizens and the infrastructures that support them.

We now can consider the whole of the model, the security stack as depicted in Figure 1. Like the OSI model from which it draws inspiration and purpose, the security stack does not define inviolable boundaries or gates, nor does it imply that one layer of security cannot exist without another. We largely have that today where critical cyber systems rely almost exclusively on Layer 2 security – the so called “bolt-on” security. We have as well the rising tide of loss to inform us that bolt-on security is an insufficient means. So the security stack aids in understanding that the complexity of the security problem requires all four layers, that the interfaces between layers are gradients, and that the layers need to exchange services in an integrated fashion. Only by addressing all four layers can we rebalance the benefit and risk of cyber systems, better understand costs to achieve the right balance (prevention) and costs if we do not (impact), and start including gradients of trust in the transactions of cyber systems. It is entirely possible that the dynamic assessment of trust can be based on an additive or aggregate value of system factors (values)

from all four layers. Layers 2 thru 4 can do little more than current point-defense functions without details from Layer 1 to inform and prioritize actions (i.e., I may need to be much more cautious at Layers 2, 3 and 4 if my Layer 1 applications and content stores – or those I am interacting with in another enterprise – are weak, but I will not know that unless I can access some details on their relative security profiles).

THE SECURITY STACK – DECOMPOSED INTO ITS PARTS

Layer 1 – Assured Systems and Content (AS&C): This layer is the set of information-communications technologies (ICT) architected and designed to operate securely within an appropriate cyber-threat environment. For example, a system designed for government information processing would be expected to operate within a higher cyber threat environment than, say, a system designed for consumer entertainment. Accordingly, a greater degree of inherent security should be applied in developing the government system. For instance, its software code should be developed using the disciplines of software assurance.

We need better intelligence regarding what is going on inside the network perimeter and what is taking place outside the network, beyond immediate control. This, in essence, is situational awareness.

Layer 1 employs technologies or methods such as data encryption or use of software assurance methodologies. Another example is whitelisting – that is, permitting only specifically authorized software to run on a system. Whitelisting involves technologies and methods to trace the root of software code back to a legitimate source. A trust anchor, as with the trusted platform module (TPM), is a fourth such example. Determining a root of trust or provenance is critical for Layer 1, and a disciplined method for configuration management is essential. Information generated in Layer 1, such as logs, can be passed to Layer 2 to establish patterns of legitimate (normal) behavior and can be used by Layer 2 to distinguish normal traffic patterns from anomalies (such as unauthorized data leaks).

Another central concept for this layer is the use of standards – engaging the knowledge of bodies such as the Trusted Computing Group (TCG), IEEE, OASIS and others to achieve rigor in the processes for assured systems and content. The ITU/T X.805 recommendation for network design is another standard, as is ISO 27000 at the policy and process level. Together, these standards provide the foundational knowledge and guidance for designing secure components, code and ICT services.

The information exchange between Layers 1 and 2 (as in the example about logs from Layer 1 being passed to Layer 2) can be extensive and requires that information go from machine to machine without human intervention to achieve speed in detecting anomalous behavior. This necessitates data structures or formats, such as XML, that facilitate the flow of information. Information flow may include state-of-health reporting (up-stack to Layer 2) to a manager (an application that manages security information), where a series of decision rules can be applied (down-stack to Layer 1), including automatic adjustments to defenses (for example, closing a port or denying an unauthorized configuration change), or alerting security analysts in an operations center. Security information exchanges can also include adjudications about levels of trust and reputation.

National critical infrastructures such as tele-communications networks, the power grid, and air travel must be protected in the interest of national security.

A variety of different professional security services enable the capabilities for Layer 1. A partial list would include cyber forensics to assess the security integrity of the software code, Common Criteria evaluations, system accreditation and certification, and security architecture and design. At the platform level, efforts focused at this layer would include use of Intel Active Management Technology (AMT), HP Systems Insight Manager and Intelligent Platform Management Interface (IPMI). For software assurance, along with reference image-based whitelisting, there is the National Institute of Standards Software Assurance Metrics and Tools Evaluation (NIST SAMATE), and the Trusted Computing Group–Trusted Middleware Suite (TCG-TMS).

Layer 2 – Integrated Security Overlay: Layer 2 is the traditional “security” layer as we know it today. It comprises several control planes across both the network and application layers. There are many forms of overlay in this layer, ranging from engineered-for-purpose hardware to software evaluation tools. Typically this is where we add defense in depth, based upon sensitivity to risk. For various reasons the security industry has evolved in a series of so-called “point solutions,” each vendor’s solution independently addressing problems at specific points in the architecture. For instance, Web application firewalls were developed to address the fundamental issues associated with Web servers facing a general purpose “anonymous” network that provides information to unknown consumers. Anti-virus software updates were built as a means to inoculate a workstation or server against known and later unknown forms of malicious software that could be downloaded by or pushed to these platforms.

Because the interconnectivity of our systems is so complex, these connection points make an enticing hunting ground for those who wish to exploit them for profit or to do harm. A large number of solutions and competitive technologies focus on providing security to these points in our networks, but information exchange among these security elements is of key importance, and they are confounded by a lack of interoperability (as in incompatible data formats from different sensors) that ultimately slow the process of correlating information needed in detection efforts. Efforts are underway among standards bodies like IEEE and TCG to address these issues; but standards take time to develop and sometimes create problems for the manufacturers of these very same solutions. Many times, additional interface technologies need to be developed.

Interoperability data formats like XACML are also being developed, and solution standards such as Trusted Computing Group Interface–Metadata Access Point (TCG IF-MAP) are being considered. Once common interfaces and object definitions are adopted, interoperability will become a more achievable goal.

Layer 3 – Intelligence: The anonymity of the Internet and certain shortcomings of TCP/IP make it difficult to learn about those who would do harm. This is the problem of attribution. We need better intelligence regarding what is going on inside the network perimeter and what is taking place outside the network, beyond immediate control. This, in essence, is situational awareness. One definition of situational awareness is

If a national level threat can manifest itself in near-zero time, an effective response requires commensurate speed.

maintaining a constant vigil over important information, but in cyberspace we can do this only to a limited extent.

We have, however, begun to address such limitations. For instance, efforts are underway to extrapolate the reputation of IP addresses on the basis of activity over a period of time, as well as “now.” Reputation services are now being considered for browser-based activities as well as for uniform resource locator/identifier/name (URL/URI/URN) content filtering. As we consider how reputation affects knowledge, we begin to appreciate the value of reputation in situational awareness, as the basis for decision support. The concept of provenance as one attribute of reputation (one can validate the source of the software code as from the legitimate provider) is just now being applied to the cyber world, but it is gaining importance quickly.

Situational awareness suffers from the multitude of languages used to convey information. We need communications mechanisms that allow us to combine data sources easily. Whether we need reputation information, source/destination pairings, or just confirmation that our DNS request went to a trusted, approved resolver, a common operating picture is essential. We typically call systems that provide such a picture *decision support* systems, because they help us conceptualize outcomes of using certain information in certain ways. Today, our situational awareness is incomplete because our decision-support systems are inadequate. They attempt to detect threats and recommend mitigations with only limited data sources (logs) and incomplete filtering rules. Efforts are underway, however, to improve decision-support systems.

Our current lack of automated defensive response and the concomitant exposure to emerging risk stems from not having the decision-support tools needed to not allow a transaction, an update or a requested action to occur. We need to do a better job in the future of developing trust so that we can automate our responses. Situational awareness is the first step toward automating defensive systems that will operate in “Internet time.”

Layer 4 – National Cyber Response: Layer 4 represents more recent considerations that are now expanding the domain of cybersecurity, where interests of national security intersect with the interests of the private sector. Layer 4 is distinct from other layers as it leaves the topic of networks and calls for a bridge – albeit limited – between the private and public sectors for specific functions consistent with the role of government as protector. Admittedly there is danger of overreach, but there is also danger in not having any reach. It has been established that a threat operates in Internet time. In contrast, the current means of exchanging threat information between government and critical infrastructures continues to operate in bureaucratic time. National critical infrastructures such as telecommunications networks, the power grid, and air space must be protected in the interest of national security. The needed exchanges of threat information cannot wait for bureaucratic time in these vital areas.

While Layer 4 also includes information-operations, this paper does not expand on the topic, focusing instead on such activities as exchange of threat information. Threat alerts must be dispatched in near-real time from government cyber intelligence organizations to the security operations

centers (SOCs) of private companies that operate targeted critical infrastructure. By way of example, Layer 4 activity is required when government authorities learn that a nation-state has employed an advanced persistent threat (APT) to steal information on power grid operations. This activity may involve not only an alert, but also close collaboration between the energy industry and the government (a) to counter the potential of the specific threat to manifest itself, and (b) to defend national interests with government action. This is not news, merely a statement of reality.

Organizations such as the Computer Emergency Response Teams (CERTs) exercise a role in Layer 4 and provide one of the few procedural standards for exchange of information. National CERTs are part of Layer 4 because they establish agreements and encourage collaborative exchange of information about threats between the private and public sectors. These are helpful beginnings, but Layer 4 is also about expanding these capabilities to respond to threats faster – as already stated, in near-real time – and about further engagement of national cyber interests in the political process. Some of these efforts involve breaking new ground and have implications

for national policy and international discourse.

The primary information exchanges in Layer 4 relate to Layer 3, where security monitoring and management activities of a private-sector SOC can receive alerts from a national-level SOC in time to act. This usually requires a person-to-person form of exchange, but there is ample room to explore automated exchanges such as bilateral situational awareness. As with the other layers, speed is of the essence in this exchange. Speed in exchanging information can be achieved only with standardization, which allows machines to handle the collecting, filtering, compiling, and exchanging of processed, decision-level information to the analysts.

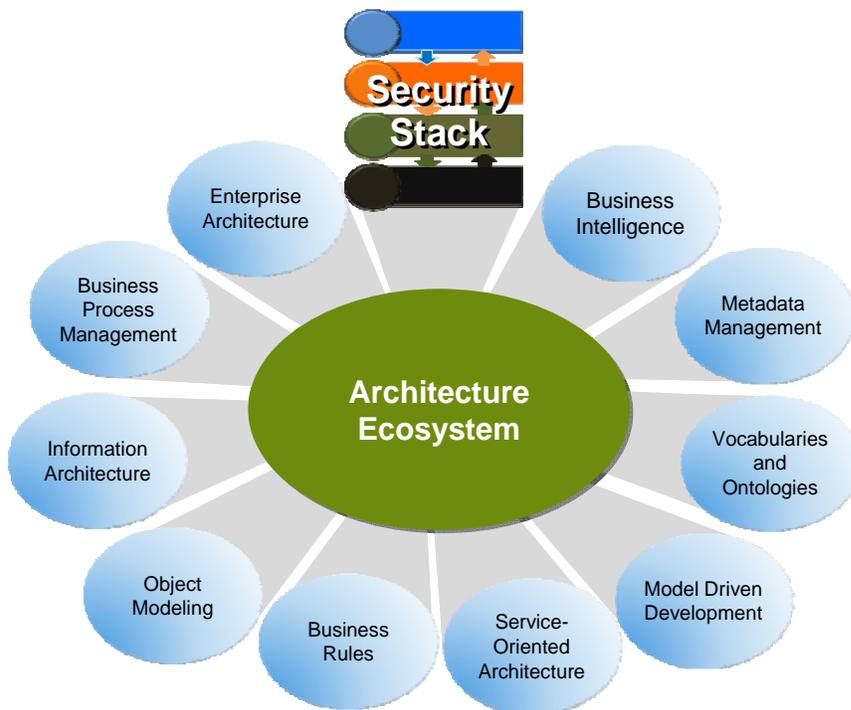


Figure 2: Architecture Ecosystem

RELATIONSHIP OF THE SECURITY STACK TO OTHER REFERENCE ARCHITECTURES

The Security Stack as part of an Architecture Ecosystem: As stated in the introduction, it is an accepted and fundamental tenet of secure environments that “security should be built in, not bolted on.” This is easy to say but not necessarily easy to do. Key to “building in” security is

Possibly, the ICT industry has nested too deeply in the forest of problems, and the security stack may provide the necessary elevation to view the full landscape.

recognition that a security stack is really one particular view of a system, system of systems, or environment. For example, a data architect views a system as a set of entities, data flows, attributes, and associations. An enterprise architect views it as a set of components, interfaces, data and stores. A technical architect thinks of the same system as a set of server nodes, firewalls, routers, disk farms, and so on. None of these views individually is *the* correct view, but each is necessary to describe an environment completely. It is time to consider the security stack as another architectural view critical for a complete information ecosystem – one that considers the purpose of use, encompasses the operating environment, and ultimately makes appropriate security design decisions by having a four-layered view. Security should not be treated as an uncoordinated installation of security point solutions.

Figure 2 adds the security stack to other architectural views. It is a part of an architecture ecosystem – a collection of architectural views (rules, enterprise architecture, data, metadata and now *security*) that collectively specify all the elements of a system and its environment. As is true of a biological ecosystem, all elements of an architecture ecosystem are interconnected and balanced. That means the security stack elements described above affect architectural elements of other views, and the elements of the other views affect the security stack elements. This interdependence is one means of assuring that security is built in and not bolted on.

This same consideration applies to the system development lifecycle. Security in its whole (as in the security stack) must be considered from inception (including capturing business function requirements), to the development of business and technical designs, and through the stages of Build, Validate and Deploy.

GENERAL BENEFITS OF THE SECURITY STACK MODEL

Many benefits derive from the uses of models, and the security stack offers the standard ones (helping to guide, explain and organize a complex set of functions). But all benefits are not created equal, and arguably, the most important benefit of the security stack is the understanding that no single layer is fully effective on its own. It is a reminder that, possibly, the ICT industry has nested too deeply in the forest of problems, and that the security stack can provide the necessary elevation to view the full landscape.

Software designed to be effective and free of vulnerabilities cannot on its own be resistant to compromise. To be fully effective within the context of the security stack, the software must be part of a well-designed network architecture that protects the software's interfaces as it exchanges protocol or end-user content information, perhaps by encrypting communication paths between network elements. Further, the security stack also reminds us that protecting software interfaces is necessary but alone not sufficient. There also must be the robust, integrated security overlay of Layer 2, and Layer 3, which would include a security incident/event manager to connect

the dots of a multi-layered, threaded, sophisticated attack. Layer 3 would also correlate (as in “security event correlation”) an external view of the threat. When appropriate, as in a national security event, Layer 4 would integrate a response with appropriate government organizations.

The security stack also can be an aid in developing policies that address issues from the network level up to national levels. Its degree of abstraction helps to drive home the point that integration is necessary, and that policies and standards must be adopted to support faster information exchange. Layer 3 of the security stack offers a case in point: Situational awareness is paramount, but it cannot be achieved fast enough to counter the threat without well-structured, integrated formats of security information exchange. So the security stack can create forcing-pressures to make collaboration (i.e., for threat information) happen not just through policy, but also at the technical protocol level where standards are most useful.

SUMMARY

Our ICT systems are inherently complex, and the degree to which they are interconnected and interdependent with critical national infrastructures, including physical systems, requires that the ICT community address this complexity. Adequate protection of ICT systems demands an adequate level of sophistication in securing them. This means that ICT systems and their content must incorporate security in their architecture and design; that the systems and the content need dedicated security activities; that intelligence capabilities must look both inside and outside the enclave perimeters and provide advance warning of threats; and lastly, that the interdependency between the private and public sector must be understood and addressed in reality – not just in talk.

Sections II and III provide the definitions and describe the role of security architectures and design in the context of other architectures, arguing for a disciplined approach aided by this model and its associated standards. The stacks disassemble the problem and provide a visualization of the necessary integration between discrete layers. The absence of this interlocking approach has put us in our current situation, with a history of known vulnerabilities coded in past software development and brought unwittingly into the present, which go undetected in the system development lifecycle. These vulnerabilities are exploited at will to do harm, with consequences that strike at the core of business and government functions. The evidence of this truth is everywhere to observe.

AUTHOR BIOGRAPHIES

Carlos Solari: Mr. Solari serves as vice president of Cyber Technology and Services for CSC. He manages the development of a wide range of cyber solutions, technologies and services related to computer network operations. Among his responsibilities is the overarching management of the development of cyber solutions and services for CSC's customer-facing organizations throughout the public and private sectors. Mr. Solari has more than 30 years of experience in information technology (IT) most recently joining CSC from Alcatel-Lucent and previously served as the Chief Information Officer at the White House. Earlier, he managed several large-scope, full life-cycle IT programs for the Federal Bureau of Investigation. In 2004, Mr. Solari was selected as a Top 100 Federal Executive by Federal Computer Week. He co-authored a book on cybersecurity entitled, *Security in a Web 2.0+ World* (2009). Mr. Solari holds a Bachelor's degree in Biology from Washington and Lee University, and a Master's degree in Systems Technologies, Joint Command, Control and Communications with an emphasis in Computer Science from the Naval Post-Graduate School, Monterey, Calif.

Dean Weber: Mr. Weber is a director and cyber solutions architect at CSC, where he provides vision and guidance for solution development within the company's Cyber Security Laboratories. With more than 30 years of experience in information and physical security, he joined CSC after serving as Chief Technology Officer at Applied Identity, which recently was sold to Citrix. Earlier, he was Chief Security Architect at Teros, a leading manufacturer of application security gateways, also acquired by Citrix. He was responsible for developing and implementing solution deployments including assessment and intelligence gathering at TruSecure/ICSA Labs (now Verizon Business Security Solutions). Mr. Weber helped found a large Midwestern reseller-integrator specializing in secure architectural design and deployment for both public- and private-sector clients, and he served for many years as its technical vice president. Additionally, he spent several years in the U.S. Navy working in physical and electronic security. Mr. Weber is a frequent speaker at information security events such as InfoWorld, ITEC, InfoSec Europe, InfraGard, Secret Service Security Roundtable, ISSA, and various focus engagements.

Victor Harrison: Mr. Harrison leads CSC's public sector Distinguished Engineering Group, providing leadership to CSC's engineering community. A widely acknowledged SOA expert, he is a member of the Board of Directors of the Object Management Group and the SOA Consortium. His articles have been published in magazines ranging from CIO Magazine to

JOOP (Journal of Object Oriented Programming), and he has spoken at industry events including Gartner's Architecture and BPM Conferences and the recently completed IT Security Automation Conference. During his career of more than 35 years, he has served as chief technical officer for a manufacturing company, has been research and development director for an ERP company, and has provided architecture and engineering leadership to over a hundred different organizations ranging from commercial enterprises to federal and state agencies to software vendors. Mr. Harrison holds a patent for a pleomorphic archetype for systems engineering. His current professional interests include modeled correctness of distributed, event-driven, and non-deterministic systems; cybersecurity; event-driven SOA; and metadata-driven architectures utilizing dynamic ontologies.



Worldwide CSC Headquarters

The Americas

3170 Fairview Park Drive
Falls Church, Virginia 22042
United States
+1.703.876.1000

Europe, Middle East, Africa

Royal Pavilion
Wellesley Road
Aldershot, Hampshire GU11 1PZ
United Kingdom
+44(0)1252.534000

Australia

26 Talavera Road
Macquarie Park, NSW 2113
Australia
+61(0)2.9034.3000

Asia

20 Anson Road #11-01
Twenty Anson
Singapore 079912
+65.6221.9095

About CSC

The mission of CSC is to be a global leader in providing technology-enabled business solutions and services.

With the broadest range of capabilities, CSC offers clients the solutions they need to manage complexity, focus on core businesses, collaborate with partners and clients and improve operations.

CSC makes a special point of understanding its clients and provides experts with real-world experience to work with them. CSC is vendor independent, delivering solutions that best meet each client's unique requirements.

For more than 50 years, clients in industries and governments worldwide have trusted CSC with their business process and information systems outsourcing, systems integration and consulting needs.

The company trades on the New York Stock Exchange under the symbol "CSC."