Aitè PARTNER. ADVISOR. CATALYST.

Black Core Rising: Microsegmentation in the Software-Defined Perimeter

DECEMBER 2019

Prepared for:

UNISYS | Securing Your Tomorrow®

TABLE OF CONTENTS

LIST OF FIGURES

EXECUTIVE SUMMARY

Black Core Rising: Microsegmentation in the Software-Defined Perimeter, commissioned by Unisys and produced by Aite Group, demystifies what the software-defined perimeter (SDP) is, how it's implemented, and, specifically, how Unisys Stealth applies software-defined microsegmentation and cloaking to prevent post-exploitation pivoting in a breach.

Key takeaways from the study include the following:

- With so many breaches over the past two decades, we aren't short of empirical data regarding the importance of organizations moving away from flat networks. Organizations must implement segmentation and microsegmentation to limit the dwell time and pivoting potential of an adversary in a breach. SDP gives organizations a more efficient and effective way of implementing microsegmentation using software.

- Segmentation and microsegmentation are different concepts. Segmentation is taking a flat network where all hosts are reachable from the other hosts on the network and breaking it into smaller subnetworks (subnets). Traffic filtering is implemented using VLAN access control lists (VACLs) or firewall rules to prevent hosts in one subnet from reaching every port and protocol to hosts in other subnets. The VACLs or firewall rules determine which subnets can talk to which other subnets. Microsegmentation brings that traffic control down even further to hosts within the same subnet. VACLs and firewall rules can take action only on traffic between subnets, not on hosts within the same subnet.

- An SDP is often conflated with the concept of software-defined networking (SDN)—the two are different but not mutually exclusive. The concept of SDN is to decouple network management (the control plane) from the flow of traffic (the data plane) in switches and routers that enable network administrators to build more intelligent, programmable networks using a central pane of glass. Before SDN, a hardware controller that manages the network, along with routers and switches, combined hardware devices with software to perform these functions. By virtualizing the network, networks can be spun up and torn down and can grow and shrink dynamically as needed. Networks can be purpose-built for the needs of specific protocols or optimized to meet the needs of specific protocols, such as H.323. SDN is now extending outside of the data center, where companies are using it to create SDNs in the wide-area network between physical locations (SD-WAN), aggregating the different types of network connections. Because the two are not mutually exclusive, they can exist together, meaning companies are even applying SDP to SDN, where the different SD-WAN connections are microsegmented to prevent pivoting between those connections. In SDN, what were previously expensive hardware-based appliances are now being replaced by software running on commodity hardware (network function virtualization). In this new hybrid and multicloud world, organizations are now using SDN to connect physical locations to their cloud service providers using microsegmentation. What this essentially means is that SDN brings the concept of virtualization to networking.

- Organizations looking to implement a zero-trust (ZT) architecture in their environment should implement SDP and consider Unisys Stealth (Stealth) as the technology to apply it. A Stealth-aware network would severely limit an attacker's ability to pivot within the network, footprint the network, reach high-value targets, harvest credentials, and perform man-in-the-middle attacks. Stealth effectively constrains traffic onto a single system, limiting the ability to move laterally around the network in order to harvest data and compromise other accounts.

INTRODUCTION

The SDP is a framework created by the Cloud Security Alliance (CSA) that applies the concept of ZT security to networking. ZT in networking describes the idea that all assets should first be authenticated and authorized before they can initiate communication with another asset. SDP implements microsegmentation using software to create trust relationships between assets so long as they are authorized to by policy.

This white paper demystifies the concept of SDP, the etymology of the term, and how it's uniquely applied by Unisys Stealth(core), deconstructing it down to its unique features and capabilities so that its efficacy, feasibility, and operational readiness for production can be measured prior to a go-forward decision by chief information security officers, security engineers, and other cybersecurity leaders looking to implement SDP technology into their environment.

METHODOLOGY

Primary research was performed for the material in this report in on-site interviews, product training, and use at the Unisys corporate office. Information on the global information grid was curated from both public-sector and private-sector data, as was information contained in unclassified material from the Defense Information Systems Agency (DISA).

ANTECEDENTS IN SDP

The concept of SDP, like many other cybersecurity concepts, has its antecedent in U.S. military applications. SDP's lineage extends back to the DISA Global Information Grid in a secure network architecture referred to as a "black core," in which both classified and unclassified data between all networked nodes is encrypted. This prevents unclassified nodes from accessing classified data.

The idea of the black core as proposed by the Department of Defense's (DoD) CIO's office was that all of the IT assets, across all of the unified combatant commands, in and out of theater, would be able to communicate cohesively while maintaining mobility, security, and survivability in information assurance. Nodes would be permitted to initiate communication with one another, within or outside of its core network, so long as they had the appropriate encryption keys to create IP Security (IPSec)-encrypted tunnels. Unencrypted user networks could communicate with nodes in the black core so long as they were permitted to by policy, where the egress unencrypted traffic would be encrypted by IPSec or High Assurance Internet Protocol Encryptor (HAIPE) gateways between the different data planes. Nodes would be authenticated and authorized before being allowed to even see other hosts in an "identity-centric approach" that relies on a need-to-know architecture, effectively blackening out infrastructure assets that aren't part of the same allowed list of assets (Figure 1).

Figure 1: Black Core in the Global Information Grid

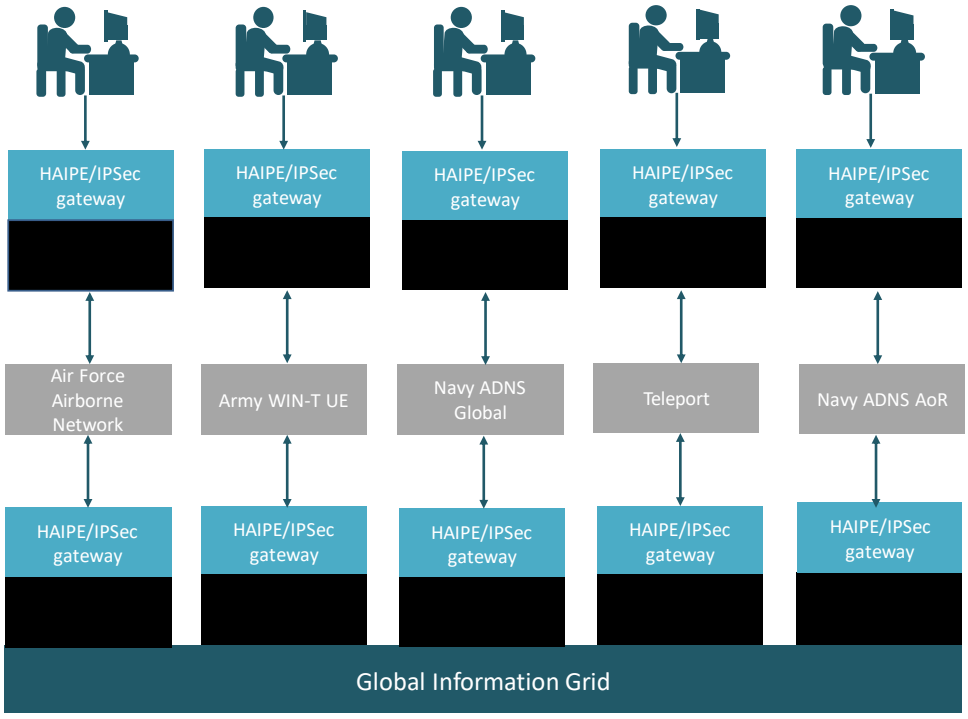

Source: Aite Group

Antithetical to the black core architecture was a new concept termed "striped core," which was proposed by Julie Tarr and Tony DeSimone from the Johns Hopkins University Applied Physics Lab in a seminal paper titled "Defining the Gig Core."[1] The paper addressed perceived challenges in applying computer network defense security controls on a black core network where 100% of the traffic between nodes is encrypted. The challenge perceived by Tarr and DeSimone in the black core design was applying pattern matching on payloads or quality of service prioritization on high-priority traffic, such as latency-sensitive H.323 and Voice over Internet Protocol (VoIP) traffic that was encrypted (Figure 2).

The striped core concept would permit decrypted traffic between nodes inside the same core, but the traffic would be encrypted when egressed outside the core to another node on another DoD component's core.[2] However, this created its own problems, whereupon if a node within the same core was breached, there would be an impact to confidentiality and integrity of the data transmitted within that core since it was being transmitted unencrypted. This also enabled pivoting and increased dwell times of adversaries who breached nodes within the same core initiated from breached assets.[3]

Figure 2: Striped Core Architecture in the Global Information Grid

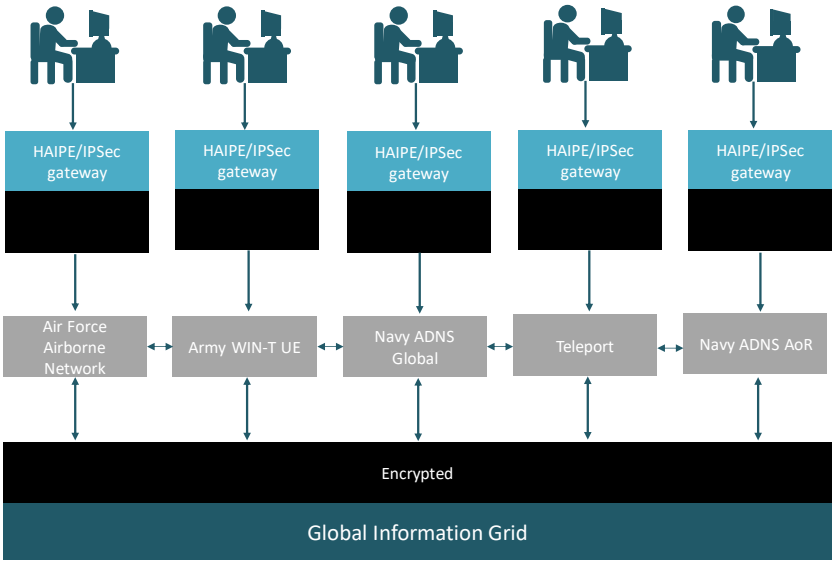

Source: Aite Group

1. Julie Tarr and Tony DeSimone, "Defining the Gig Core," IEEE, 2007, accessed December 4, 2019, https://ieeexplore.ieee.org/document/4455016I.
2. DoD CIO, "Department of Defense Global Information Grid Architectural Vision," DoD CIO, June 2007, accessed December 5, s, http://www.acqnotes.com/Attachments/DoD%20GIG%20Architectural%20Vision,%20June%2007.pdf.
3. David F. Carr, "Building a Protective Black Core for the Global Information Grid," September 9, 2009, accessed December 4, 2019, https://defensesystems.com/articles/2009/09/02/cyber-defense-black-core.aspx.

THE BEGINNING AND END OF THE VPN ERA

Following the development of the IPSec protocol suite in 1995, Microsoft unveiled a virtual private network (VPN) built on a new protocol it termed point-to-point tunneling protocol (PPTP), providing the capability for hosts to create private networks with one another in private, encrypted tunnels that allowed them to communicate over inherently untrusted networks such as the internet. Other types of tunneling protocols would later surface, including L2TP/IPSec and other protocols for the secure exchange of keys, such as Internet Key Exchange (IKE).

However, VPN technology has not changed over the last couple decades, and many believe that—as digital transformation pushes what were once on-premises servers into the cloud—the need for a "classic" VPN has been diminishing quickly. Also compounding the issue is an increasing movement toward virtualization and migrations away from monolithic applications to microservices containers. Companies are now dismantling their edge VPN concentrators as resources needed by remote employees are moved into the cloud.[4]

The dissolution of the perimeter has not been the only threat to the VPN market. The number of breaches perpetuated by site-to-site VPNs between companies and their suppliers through supply-chain hijacking has continued to rise, which has been a further black mark on "classic" VPN architectures.[5]

4. "A Brief History of VPN," Golden Frog, GmbH, June 22, 2016, accessed December 11, 2019, https://www.goldenfrog.com/blog/brief-history-of-vpns.

5. Patrick Sullivan, "The Death of the VPN - It's Time to Say Goodbye," SC Magazine, March 21, 2019, accessed December 5, 2019, https://www.scmagazine.com/home/opinion/the-death-of-the-vpn-its-time-to-say-goodbye/.

RISE OF THE SDP

To fully explain SDP, we need to break it down into its smaller parts, describing the distinctive differences between SDN and SDPs, the concept of "need to know," what microsegmentation is versus segmentation, and how IPSec and ZT are inherent to SDP.

An SDP is often conflated with the concept of SDN—however, the two concepts are different. The concept of SDN is to decouple network control (the control plane) from forwarding functions (the data plane) in switches and routers to enable network administrators to build more intelligent, programmable networks using a central pane of glass. SDN allows administrators to dynamically respond to changing network demands of the systems and applications, using it at the software layer rather than laborious manual configuration changes at individual hardware devices, as the network evolves over time; is done at a centralized, high level; and is more vendor-agnostic than in years past. This reduces the capital expenditure of expensive hardware and operational costs due to lost time and productivity. SDN virtualizes much of what used to be done with expensive hardware devices, such as the virtualization of physical servers, using what's found in hypervisors like VMWare and Hyper-V.

Thus, it would be accurate to consider SDN as a centrally programmatical way to configure network elements to implement connectivity between systems, whereas SDP is a way to configure systems to control their communications with each other without regard to the network connectivity between them. This segmentation used to be done with hardware at the switch level using VACLs or firewall rules.[6]

With SDP, administrators can isolate machines into their own isolated communities (such as an in-scope PCI environment) within the much larger corporate network. SDP can even bring remote users into that same private community, appearing as local machines in the corporate office, despite the geographical disparity between the hosts (Figure 3).

6. "Software-Defined Networking (SDN) Definition," Open Networking Foundation, accessed December 5, 2019, https://www.opennetworking.org/sdn-definition/.

Figure 3: Example of Microsegmentation in SDP

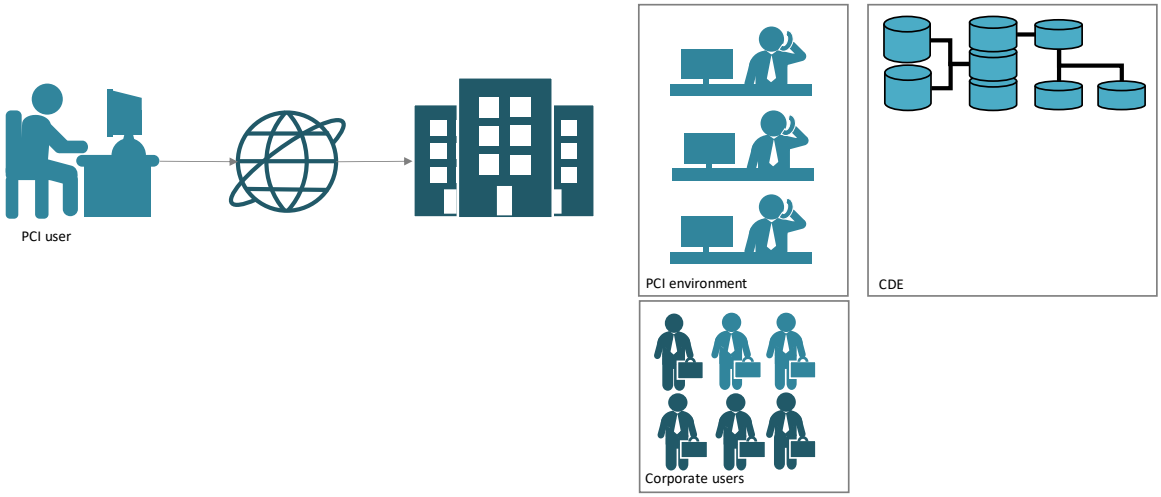

Source: Aite Group

NEED TO KNOW

The concept of "need to know" in military argot purports that just because you have a clearance level, you don't necessarily have a "need to know" information that is classified at that same level. While a clearance at that level means you can have access to it, it doesn't mean that the data owners necessarily have to grant you access to it simply because you're cleared to see it. You must have a "need to know" the information. This same concept applies to SDP. Simply because your host is a member of the corporate network, it doesn't necessarily mean you should have a need to know all data that's transmitted and stored within it. For example, just because you're an employee of a company, doesn't necessarily mean you have a need to view the internal employee investigation records in the human resources department.

In SDP, an asset must be granted permission to speak to another asset through policy definition, which is enforced by the SDP controller. Authentication and authorization of hosts are done differently, depending on the SDP solution, but they typically use the concept of SDP endpoints and SDP controllers. In SDP solutions that apply encryption between endpoints, when an SDP-enabled endpoint wants to initiate a session with another, it must have the proper encryption keys to initiate that session, which is again governed by policy at the SDP controller.

MICROSEGMENTATION VERSUS SEGMENTATION

The term "segmentation" should be viewed as more of a macro-view of segmenting the network, achieved by dividing the network into multiple broadcast domains, separated by routers and firewalls. VLANs were employed to extend the concept further by making the separation of assets more logical than physical. In classic network segmentation, segmentation was implemented at the network elements that glued systems together, rather than in the endpoints themselves, operating within a purely logical policy framework as applied in microsegmentation. Microsegmentation should be considered a more surgical precision in the

carving out of what hosts can talk to what hosts. Within the larger network, this is what's referred to as east-west traffic. Simply put, microsegmentation is the compartmentalization of networked assets into groups permitted to talk to each other over secure, encrypted tunnels.

Segmenting a network is typically an approach in which network administrators create individual VLANs based on the role of the asset—such as a user VLAN where employees are placed and a facilities VLAN where internet of things (IoT) devices such as badge readers and CCTV cameras are placed—with each VLAN typically being a different subnetwork (subnet). Network segmentation can be implemented to apply filtering rules on network traffic, whether it is north-south or east-west.

In network segmentation, assets are placed into a different IP space, and VACLs or firewall rules are used to govern what ports and protocols these assets can talk to when initiating communication outside of their network segment. These are considered IP-based rules. But microsegmentation doesn't care about the IP address of assets—rather, the user's identity and their "need to know" as defined by policy. In order to enforce access policies within these microsegments, a new approach was needed, which is met through SDP using a purely logical policy framework where the endpoints themselves enforce access policies instead of by network elements.[7]

IPSEC

IPSec is an entire suite of protocols designed to provide confidentiality, integrity, and authenticity of data transmitted between systems. The receiving system can be certain that all the packets it receives were sent by the system that initiated the session, that none of the data was changed in transit, and that the data cannot be used by any intermediary.

IPSec is used by various traditional architectures (such as VPNs) in north-south traffic flows, but SDP uses IPSec to create discrete and authorized communication sessions between hosts, as governed by the SDP policies, in east-west directions.

It's important to note that encryption in SDP is purely optional and can use VPN tunnels (including IPSec tunnels) to provide greater security in north-south as well as east-west traffic. Communications through IPSec tunnels cannot be observed, intercepted, altered, or impersonated by any intermediary system.

ZERO-TRUST SECURITY

The concept of ZT security was created by John Kindervag in 2010, a principal analyst at Forrester Research. ZT describes the concept that organizations should not trust users, systems, or data inside or outside its edge. Trust nothing.

7. Katherine Teitler, "How Microsegmentation Differs From Network Segmentation," Edgewise Networks, March 12, 2019, accessed December 5, 2019, https://www.edgewise.net/blog/how-microsegmentation-differs-from-network-segmentation.

This is the very tenet of SDP—authentication and authorization of assets based on identity and whether they have a need to access the system or data being requested, controlled by policies, which are based on business requirements.

UNISYS STEALTH

Unisys Stealth is an on-premises and cloud-deployable solution for implementing SDP using only software. Built on the concept of communities of interest (COIs), Stealth-enabled assets are able to communicate with one another so long as they have permission by policy. Non-Stealth-enabled assets are unable to see the data being transmitted between Stealth-enabled assets, because the data is encrypted in IPSec tunnels.

COI membership is dynamically negotiated between Stealth-enabled endpoints when the initiator presents its COIs to an endpoint it wishes to initiate a connection to. The target can determine whether the two endpoints have a COI in common, but without revealing any other of the initiator's COIs to the target and without revealing any of the initiator's COIs to someone sniffing the network. If the initiator does not present any acceptable COIs to the target, the target does nothing. Therefore, the target appears "dark" to unauthorized traffic, as if the host doesn't exist. If one of the initiator's COIs is acceptable, then the target presents the common COI back to the initiator in a way that allows the initiator to verify that the target indeed possesses the same COI. This allows Stealth to dispense with the controller element in the SDP architecture.

With Stealth, organizations that used to spend hours laboring over switch configurations, rerunning network cables, and making other hardware-based network changes can now simply push Stealth agents out to assets to "hide" them in predefined communities, without the need to make changes at the infrastructure level. Assets that are running network operating systems not supported by the Stealth software agent can still be brought into COIs in an agentless mode, additionally giving organizations the ability to Stealth-enable critical assets, such as network equipment or proprietary devices that include IoT, supervisory control and data acquisition (SCADA) equipment, and legacy computing technology. Secure virtual gateways used in these deployment types (referred to in the Stealth environment as clear-text workstations, which are unable to run the Stealth software agent) allow them to still participate in COIs.

COMPONENTS

The Stealth architecture has several components. An Enterprise Manager instance provides a web-driven user interface where Stealth policies are defined, licenses are managed, logs are stored, and the configuration and administration of the Stealth environment is performed. Enterprise Manager should be considered the nerve center of the entire Stealth infrastructure, but it does not participate directly in Stealth policy enforcement. Authorization Service instances authenticate each endpoint identity, refer to policy definitions to determine the appropriate policies for the identity, and provide the policies to each endpoint securely.

COMMUNITIES OF INTEREST

Stealth COIs form the basis of network access policies in the Stealth architecture. COIs can be created around types of data classification, such as secret, or top secret, or sensitive earnings reports for publicly traded companies. COIs also can be created around specific network segments, such as in-scope PCI assets.

Only users in the same COI are allowed to communicate with one another, and users can be a member of more than one COI. Every COI has its own unique encryption key shared among COI members. Therefore, in order for an asset to be able to communicate with another asset, it must have the correct key for that COI, which is why it's imperative that Stealth endpoints that you wish to be able to communicate with one another are assigned to the same COI.

Each COI is represented by a unique cryptographic key. Endpoint identities are authenticated by Stealth against a trusted authority, such as a certificate authority (CA) or Lightweight Directory Access Protocol (LDAP) directory. Each identity is mapped to a role, according to policies defined in the Enterprise Manager. Each role has network access policies defined in the Enterprise Manager. The network access policies for each role are primarily a set of COI keys. Each COI in a role may have rules that further constrain how the role may process network traffic within that COI. Each role may have Clear Text exception rules, which define non-COI network traffic that the role may process.

CONCLUSION

- The decision for cybersecurity leadership should not be whether to implement microsegmentation or segmentation—organizations should be doing both. The network should be segmented based on the role or data classification transmitted, processed, or stored by the asset, and identity-based authentication and authorization should be used to determine which assets should be allowed to communicate.
- SDP can be an effective network security control to make an organization more resilient to traditional cybersecurity attacks, such as Denial of Service (DoS) attacks, credential stuffing, man-in-the-middle attacks, and more. Additionally, policies within Stealth can be rapidly reconfigured so that compromised assets are isolated in an effective incident response posture.
- When deciding to implement SDP, organizations should ensure all of the proper stakeholders from outside cybersecurity are involved, such as system owners, data owners, and infrastructure and operations.
- To implement a true ZT architecture in the enterprise, SDP should be considered a necessary first step in ensuring that assets are authenticated based on their identity, and their specific network access is authorized according to policies based on business requirements.
- SDP can be an effective security control for controlling network access of operational technology, such as SCADA and industrial control systems, decoupling them from the corporate network. Physical security controls, such as CCTV cameras, should never be placed on a flat enterprise network, as they can be leveraged to pivot into the internal network of the organization unless properly isolated using SDP.
- Unisys Stealth is a multilayered solution for implementing SDP into both on-premises and cloud environments to secure workloads and implement microsegmentation across hybrid and multicloud environments, as well as to bring remote, untrusted networks, such as work-from-home users and their home networks, into trusted COIs.
- Stealth eliminates the controller, introducing Authorization Services as the elements that authenticate identities and provide policies. These enhancements make Stealth more naturally resilient and easily scalable.
- In my hands-on review of the technology, Unisys definitely “gets” microsegmentation and does it right. As a solution that works on-premises or in the cloud, across mobile devices to cloud workloads, Stealth should be the choice for anyone who is serious about designing and implementing their ZT security strategy.

ABOUT AITE GROUP

Aite Group is a global research and advisory firm delivering comprehensive, actionable advice on business, technology, and regulatory issues and their impact on the financial services industry. With expertise in banking, payments, insurance, wealth management, and the capital markets, we guide financial institutions, technology providers, and consulting firms worldwide. We partner with our clients, revealing their blind spots and delivering insights to make their businesses smarter and stronger. Visit us on the web and connect with us on Twitter and LinkedIn.

AUTHOR INFORMATION

Alissa Knight
+1.206.765.7434
aknight@aitegroup.com

For more information on research and consulting services, please contact:

Aite Group Sales
+1.617.338.6050
sales@aitegroup.com

For all press and conference inquiries, please contact:

Aite Group PR
+1.617.398.5048
pr@aitegroup.com

For all other inquiries, please contact:

info@aitegroup.com