WHITE PAPER – **AUGUST 2017**

CORE PRINCIPLES OF CYBER HYGIENE IN A WORLD OF CLOUD AND MOBILITY

vmware®

Table of Contents

Introduction

Cybersecurity is a top concern at the highest levels of government and industry worldwide. More than ever, government and corporate leaders – from Senators and Members of Parliament to CEOs and Board Directors – are deeply engaged in ensuring effective cybersecurity strategies are in place at government agencies and companies.

Yet as investments in cybersecurity accelerate, breaches continue to occur at an alarming frequency. Something is not working. What is it? And how do we fix it? There are many theories on what to do – everything from following new governance frameworks to deploying new products and services.

At VMware, we believe that more effective information security won't be achieved by following a new framework or buying a particular product. The answer is to architect security in, rather than bolting it on as an afterthought. This has been inherently difficult for organizations to achieve, but new capabilities provided by cloud and mobile computing now make it feasible, if not essential.

Moving to a more effective approach to security requires taking two fundamental steps: implement basic cyber hygiene and focus on protecting the "crown jewels"– mission-critical business applications.

In this paper, we propose five core principles of cyber hygiene as a universal baseline: the most important and basic things that organizations should be doing. The concepts are not new but are key in moving to more effective security. They are rooted in well-established frameworks such as the NIST Cybersecurity Framework (CSF) and are technology-neutral. In the most devastating data breaches over the last few years – from Target to Sony to the U.S. Office of Personnel Management (OPM) – we think effectively adhering to these principles would have made a meaningful difference.

Still, implementing core principles of cyber hygiene effectively is not easy and has eluded organizations for years. Therefore, we also propose that organizations focus security efforts on protecting applications, specifically the mission-critical business applications that are their crown jewels. This can greatly improve the effectiveness of security.

This paper is intended to help government and corporate leaders understand specific problems with current cybersecurity strategies and how to move to a better approach. It is written for leaders who are engaged in cybersecurity issues but not necessarily technical experts. For security practitioners and others who may be interested in the more technical details, we provide a set of appendices including practical suggestions for implementation.

Improving cybersecurity is high on the agenda for government and industry. As experts in cloud and mobility, we are proud to contribute our unique perspective to improving cybersecurity. We believe it's a valuable vantage point from which to tackle information security challenges. We bring our ability to see through a different lens.

"CYBER HYGIENE" DEFINED

This term has various meanings. We are using it to refer to the basic things that an organization should have in place for cyber defense.

This is different from another common view of cyber hygiene, which refers to what online consumers do to protect their personal online activities from infection.

The Cybersecurity Conundrum

Global spending on security continues to accelerate, with an estimated compound annual growth rate (CAGR) of 8.7% through 2020.[1] Yet the annual number of data breaches in the U.S. hit an all-time record high last year.[2] Corporations and governments worldwide are losing nearly $500 billion a year from data breaches.[3] Clearly something is not working. What is it? What can we do about it?

No shortage of guidance

The cybersecurity deficit certainly isn't due to a lack of guidance regarding what organizations should do to protect information. There are numerous, well-accepted government and industry standards for information security in the U.S.A. and globally, including NIST, ISO, SANS, and many more. They all point to a comprehensive list of generally agreed-upon best practices.

Complexity is overwhelming

Current approaches make it impossibly complex to implement comprehensive best practices across an enterprise IT environment. There is a huge assortment of security tools to manage: firewalls, anti-virus, intrusion prevention systems, and threat detection systems, to name a few. Each tool has an enormous number of rules to manage. Each one must be set up to enforce access control and/or information protection policies at enterprise scale–for all users and systems across the enterprise. In some cases, this could mean literally millions of rules. It's a configuration nightmare.

Change is constant

And security tools aren't just set up once and then forgotten. Throughout an enterprise, systems need to be updated constantly in order to keep up with the ebb and flow of business activity and protect against newly-discovered vulnerabilities.

Automation is out-of-reach

Although organizations have many tools in place to automate security tasks, the tools can't be used in concert in a fully automated fashion. Their functions can't be easily coordinated since they all use different ways to label the systems they are protecting.

Organizations may also be reluctant to fully automate security for fear of breaking something. For example, if they do an automated update, it might shut down a critical system. They often don't have enough information on how the patch might impact systems.

Responding to alerts is onerous

Another difficulty is the volume of work involved in following up security alerts. Each of the many security tools in an organization sends out thousands of alerts per day, in some cases, thousands per hour. Each tool has its own separate management console so the team has to be watching multiple screens. Prioritizing alerts is difficult and responding requires a lot of investigation. For example, a detection tool might indicate there is suspicious activity in the network, but provide no specifics on the affected systems, the risk level or possible actions.

Organizations rely heavily on humans to carry out security functions but the shortage of cybersecurity talent means there aren't enough to go around.

---

[1] Worldwide Semiannual Security Spending Guide, IDC, March 2017

[2] Identity Theft Resource Center (ITRC) Data Breach Report 2016

[3] Net Losses: Estimating the Global Cost of Cybercrime, Center for Strategic International Studies, June 2014

EDUCATION PROCESS

IT professionals should be committed to designing security into systems. Developers should learn a minimum amount of code-security skills. System architects should sign for security outcomes. Foundational security knowledge should be as familiar as knowledge of computing, networking or storage.

End users should be aware of the risks and their responsibilities in protecting information. Security basics should be as well understood as going to a website or checking email.

MICRO-SEGMENTATION:

Protecting the IT environment by breaking it up into smaller parts is similar to the use of compartments on a ship. It makes the ship easier to protect. If the ship is damaged in one area, the damage is contained to that area.

Two Steps to More Effective Security

Recent advances in cloud and mobile computing now make it possible to simplify and more fully automate security. There are two fundamental steps to take: implement cyber hygiene and focus on protecting the crown jewels – mission-critical business applications.

Step One: Implement core principles of cyber hygiene

These are the most important and basic things that organizations should be doing.

The Underpinning: Education

A mandatory education process should be in place for everyone from IT professionals and business leaders to employees and third-party contractors (see sidebar).

The Core Principles

With education firmly in place, these five principles are key in moving to more effective security:

1. Least Privilege	Users should be allowed only the minimum necessary access needed to perform their job and nothing more. And system components should be allowed only the minimum necessary function needed to perform their purpose and nothing more.
2. Micro-segmentation	The whole IT environment should be divided into small parts to make it more manageable to protect and to contain the damage if one part gets compromised (see sidebar).
3. Encryption	For critical business processes, all data should be encrypted, while stored or transmitted. In the event of a data breach, stealing critical files should only result in obtaining unreadable data.
4. Multi-factor Authentication	The identity of users and system components should be verified using multiple factors (not just simple passwords) and be commensurate with the risk of the requested access or function.
5. Patching	Systems should be kept up to date and consistently maintained. Any critical system that is out of date is a meaningful security risk.

Well-established principles

The core principles are not new concepts. They are rooted in well-established principles. For example, they map to various functions within the NIST CSF (see appendix 1). They are only a fraction of what the NIST CSF and other frameworks cover, however they are the key enabling principles for moving to a simpler and more automated approach.

Doing these five things, well and consistently, would make cyber-attacks much more difficult to carry out and far less damaging. Even in the most devastating data breaches over the last few years we think effectively implementing these principles would have made a meaningful difference (see below).

Major breaches where core principles not effectively implemented

Principle	Examples of breaches Note: Many factors lead to a data breach; these are examples where not effectively implementing a core principle contributed to a breach; there would also have been other factors.
1. Least Privilege	If a least privilege environment has not been effectively implemented and users are provided with higher levels of access then they need, attackers can steal these credentials (user name and password) and gain broad access to systems. For example, in the Target and Sony breaches, attackers were able to gain administrative-level privileges.
2. Micro-segmentation	If micro-segmentation has not been effectively implemented, attackers can break into one part of the network and then easily move around to other parts. For example, in the Target breach, after an initial intrusion into the HVAC system, the attackers were able to move around to the payment network system. In the Sony breach, the attackers were also able to move around from one part of the network to another. In the case of the OPM breach, the attackers obtained access to OPM's local area network and then pivoted to the Interior Department's data center.
3. Encryption	If encryption has not been effectively implemented, attackers can exfiltrate data in readable form. For example, after a data breach at Royal & Sun Alliance Insurance PLC, government investigators determined that the company had not adequately encrypted the data.
4. Multi-factor Authentication	If multi-factor authentication (MFA) is not effectively implemented, attackers can obtain passwords and use them to access systems. For example, in the OPM breach, if the contractor logons had been enforced with a risk appropriate level of MFA it would have limited the ability of the attackers to use the stolen credentials of the government contractor. In the case of the breach at LinkedIn, the hack exposed inadequately protected passwords of 100 million users. Since consumers often use passwords on multiple sites, MFA would have reduced the risk.
5. Patching	If patching is not effectively implemented, attackers can exploit open holes in systems. For example, the WannaCry ransomware exploited a known software vulnerability for which a patch was available. Organizations that fell victim had failed to effectively patch.

Not easy to implement core principles effectively

Security professionals at most organizations are very familiar with these principles. In fact, even at organizations that have experienced a breach, the security team likely tried to implement them. The problem is that it is very difficult to do given the current approach to security that most organizations take, using the tools and techniques available to them.

"Crown Jewels = Critical Applications"

Step Two: Focus on protecting *individual* critical applications

The next step is to focus on protecting *individual* critical applications. This will make it easier to effectively implement the core principles of cyber hygiene.

Focusing on critical applications puts the focus where it should be: on the crown jewels. Ultimately, an organization's crown jewels are its mission-critical business applications and the data within them. Examples include: an enterprise financial application that processes sensitive data in creating the company's financial statements; an ordering application that fulfills customer orders, including storing personal information and credit card data; an HR application that contains confidential employee data; and an R&D application that contains trade secrets. The application is the mechanism for accessing and interacting with the data.

Even though the goal of information security is to protect these crown jewels, current approaches are focused on protecting the IT infrastructure, like routers (hardware that routes traffic on a network) or servers (computers that provide processing power). Protecting the IT infrastructure is necessary but not sufficient.

Take a risk-based approach

It is the critical applications and data that are of the value to the business. The compromise of these assets represents significant risk for the organization. The infrastructure provides the things an application needs to operate but is not itself the critical asset.

Get more specific

Focusing security on the infrastructure isn't specific enough. It's like trying to protect all the houses in a community by putting a fence around them with a locked gate. It would be more effective if you focused on protecting each individual house – see diagram 1 below.

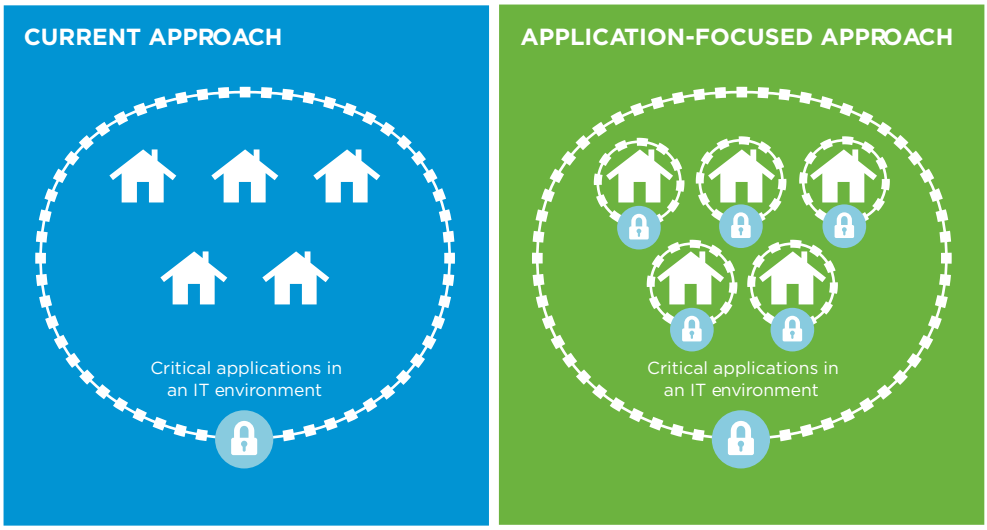

Diagram 1: The current approach to protecting an IT environment is like trying to protect the houses in a community by putting a fence around them, with a locked gate. It would be more effective to focus on the houses (critical applications) by adding fences and locks to each one.

Control access to each individual application

With current approaches, it's hard to effectively achieve security goals, such as ensuring only minimum necessary access. For example, a firewall is often set up at the perimeter of the whole enterprise (like the fence around our whole community) to control access to a group of applications, which can often be thousands of applications. Instead, there should be a firewall set up to control access to each individual critical application (like each individual house), allowing only access by the users and system components that absolutely need access to that one application (house).

Security also needs to get more efficient. Imagine that the guards at the gate get a phone call alerting them to unusual activity somewhere in the community. The guards might spend all day looking around the community looking for the unusual activity. It would be more efficient if the guards knew exactly which house to go to, if the house was empty or filled with valuables, and if the activity was normal for that house – see diagram 2 below.

Monitor with specific knowledge of the application

This is similar to information security monitoring systems. They typically send out an alert indicating an intrusion into the network or part of the network with no specifics on the application. The cybersecurity team has to spend a lot of time investigating. It would be better if the alert would indicate which application was affected, how critical it was, and if the activity detected was legitimate for that application.

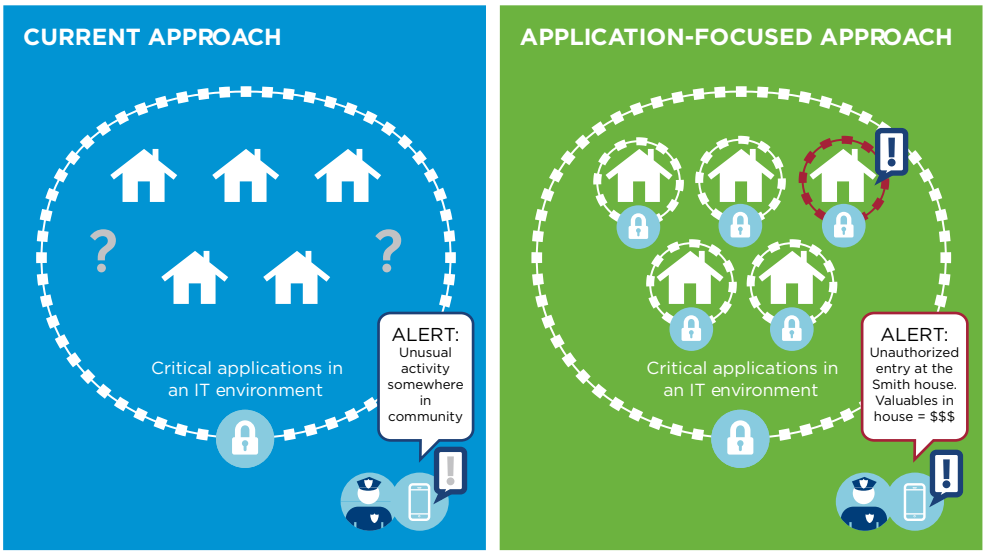

Diagram 2: The current approach to monitoring an IT environment is like alerting a guard that there is unusual activity somewhere in the community. It would be more effective if the guard knew exactly which house (critical application) and what was happening.

MODERN APPLICATIONS: DISTRIBUTED AND DYNAMIC SYSTEMS

- Each individual application is a "system"of components.
- Software functions (or services) use a pool of resources: networking, processing, memory and storage.
- Resources are spread across an IT environment (possibly across the organization's own data center and cloud providers).
- The pool of resources is shared by many applications.
- The use of resources changes rapidly over time.

Why aren't organizations already doing this?

If focusing on individual critical applications would make security more effective, why aren't organizations already doing this? Given the technologies and techniques that most organizations currently use, it's simply not feasible.

Current approaches lack the ability to delineate individual applications

Traditionally, applications were designed to have all of the application's components reside in a single, static machine. But modern applications are designed as distributed and dynamic systems. The components are spread across multiple machines, with software functions using a pool of shared resources that changes over time (see sidebar). With current approaches, security tools are not able to recognize or understand individual applications.

With current approaches, security tools:

- Cannot identify that "these components" make up "Application A"
- Don't know which users should have access to "Application A"
- Don't know which system components should be allowed to communicate with each other as part of "Application A"
- Are unable to keep track of "Application A" as it changes, for example as the software uses different hardware resources

Applications continue to evolve

Newer applications consist of smaller software functions that are even more dynamic. Therefore the current approach, centered on protecting the infrastructure, will have even less success in protecting individual applications. Therefore there is some urgency in making the shift to an application-focused approach, as the issues are only going to get worse.

Cloud and Mobile Computing Now Make it Possible

With advances in cloud (private and public) and mobile computing, organizations have the capabilities required to focus on protecting individual applications, paving the way to more effective security.

Use application-focused capabilities

Specifically, cloud and mobile computing provide the ability to:

Capability I: Recognize an individual application and establish a baseline reference for it

- Identify what components make up the application
  - Gain visibility into applications
- Know how an application is supposed to operate and how it actually operates at runtime
  - Know who and what needs access and how they interact
- Use this reference information in protecting the application
  - Refer to this information for setting up the security tools

Capability II: Compartmentalize system components into individual applications

- Group together all the system components that make up one individual application
- Associate all these system components by putting a logical boundary around the grouping
- Use the boundary to uniquely label the application

Capability III: Position defenses around each individual application

- Determine what can go in/out of the boundary around the application
- Align the security tools to the application boundary
- Set up the security tools by referring to the baseline reference information and using the label provided by the boundary
- Create an application-specific rule set for protecting one individual application
- Track the application and adjust protection as it changes

See diagram 3 below for an illustration of using these capabilities to effectively protect modern applications. For more details on the capabilities and what they enable organizations to do, see appendix 2.

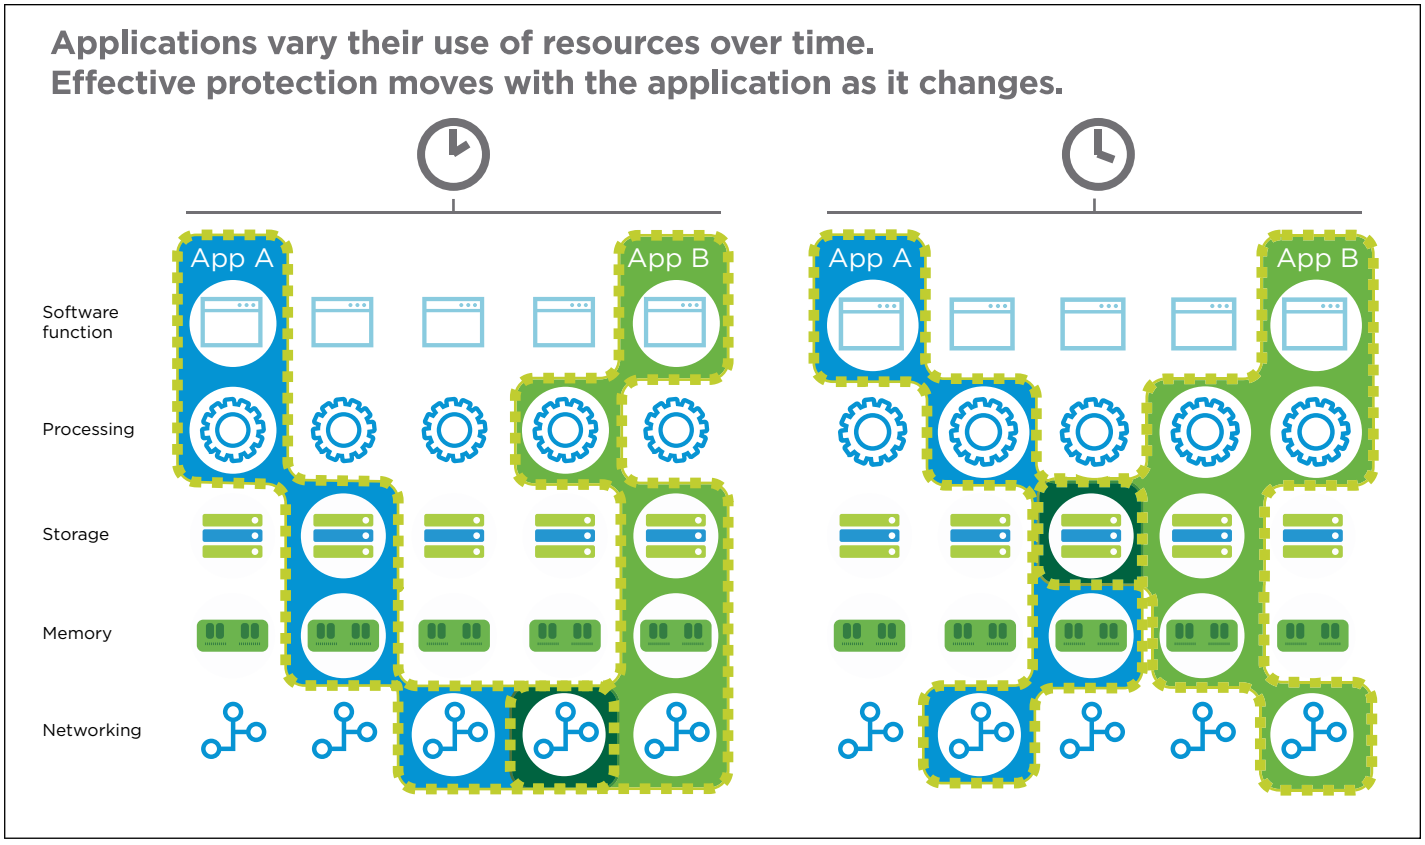

Diagram 3: A modern application is a distributed and dynamic system. It uses a pool of shared resources and varies this use over time. To protect an application effectively requires identifying all of the software and hardware components that make up the application, grouping these together, putting a boundary around them, labelling them as "Application X", and then positioning defenses around the boundary. The boundary and defenses should move with the application as it changes.

Effectively implement the core principles

If we take an application-focused approach, the core principles can then be effectively implemented. Security becomes much simpler and easier to automate:

Principle	Application-focused approach	More effective implementation.
The Underpinning: Education	A mandatory education process should be in place for everyone from IT professionals and business leaders to employees and third-party contractors, with a focus on applications.	Education will be more relevant, tailored to the applications that IT professionals and/or users work with.
1. Least Privilege	Users should be allowed only the minimum necessary access needed per individual application to perform their job and nothing more. System components should be allowed only the minimum necessary function needed per individual application to perform their purpose and nothing more.	User access and system component function will be more tightly controlled. It will be more difficult for attackers to find ways to gain access, alter processes, or hijack interactions (both "system to system" and "user to system").
2. Micro-segmentation	The whole IT environment should be divided into small parts by setting up boundaries around individual applications to make it more manageable to protect and to contain the damage if one part gets compromised	Movement within the IT environment will be significantly inhibited. If attackers do make it into one part, they will be confined to a very small part (i.e. a single application) and find it difficult to reach other parts.
3. Encryption	For critical business processes, all data should be encrypted while stored or transmitted by the components of an individual application. In the event of a data breach, stealing critical files should only result in obtaining unreadable data.	The distribution of the keys required to lock/unlock the data will be simplified since it is managed for each application individually. It will be more feasible to implement encryption comprehensively.
4. Multi-factor Authentication	The identity of users and system components should be verified using multiple factors (not just simple passwords) and be commensurate with the risk of the requested access or function for an individual application.	Enforcing a risk-appropriate level of multi- factor authentication (MFA) for every request will be more feasible since it is managed per application. Attackers will find it more difficult to perform attacks if they can no longer simply steal or guess passwords.
5. Patching	Systems are kept up to date and consistently maintained based on the knowledge of each individual application. Any critical system that is out of date is a meaningful security risk.	Patching will be much easier to do consistently, knowing which application's components are affected and the possible impact to systems. Attackers will find it much harder to find vulnerable systems to exploit.

Start by classifying applications

With an application-focused approach, the security team can zero-in on the most important assets, i.e. critical applications, rather than spreading investments thinly across the infrastructure. Organizations begin by classifying applications to ascertain criticality and prioritization, so they can put more effort into the most critical applications. Keep in mind however that all applications need some level of protection.

Improve the effectiveness of existing security tools

An application-focused approach enables organizations to make the most of their security tools:

- Reduce misconfigurations of security tools
  - Rule sets are simplified: application-specific rules are applied to each individual application
- Set up the security tools to work in concert
  - All security tools–firewalls, anti-virus, intrusion prevention systems, and threat detection systems–use the same label (the application's boundary) to identify the asset they protect
- Interpret and act on alerts more easily and quickly
  - Alerts from the security tools identify the application and provide information on the level of priority and possible courses of action
- Use security tools in a more fully automated way
  - Security tools' functions can be coordinated; the protection, monitoring, and response activities can be organized around individual applications
- Decrease the cost of security operations
  - There will be fewer alerts generated and less time spent on investigation

Architect in security

Typically, security is "bolted on." Application teams build an application, infrastructure teams build a relatively generic infrastructure capable of handling all the applications, and then the security team is asked to secure all of it. Security tools are deployed but not integrated into the fabric of applications.

An application-focused approach will require an architectural shift. It will not be achieved simply by buying a particular security appliance or upgrading software. It requires embracing the unique properties of cloud and mobility technologies and using these to security's advantage (see appendix 3 for more info).

Cloud and mobility technologies provide an overlay architecture that can be used to integrate security, not only for new applications but also existing applications. For practical implementation suggestions in the data center and for end user computing, see appendix 4 and 5.

Conclusion

By taking two fundamental steps: implementing the core principles of cyber hygiene and focusing on protecting the application – organizations can move to more effective information security. Cloud and mobility computing now make it possible and provide a way to architect in security. As IT environments continue to evolve, this updated model can help ensure that an information security program is not only more effective today but also prepared for the future.

Appendix 1: Mapping of Core Principles to the NIST CSF

The core principles of cyber hygiene are rooted in well-established principles. For example, all of them map to various functions within the NIST Cybersecurity Framework (see below). These principles are only a fraction of what the NIST CSF, and other frameworks, cover. However, they are the key enabling principles for moving to a simpler and more automated approach to security.

Core Principles	NIST CSF sub-categories
Underpinning: Education	PR.AT: The organization's personnel and partners are provided cybersecurity awareness education and are adequately trained to perform their information security-related duties and responsibilities consistent with related policies, procedures, and agreements.
1. Least Privilege	PR.AC-4: Access permissions and authorizations are managed, incorporating the principles of least privilege and separation of duties PR.PT-3: The principle of least functionality is incorporated by configuring systems to provide only essential capabilities PR.IP-1: A baseline configuration of information technology/industrial control systems is created and maintained incorporating appropriate security principles (e.g. concept of least functionality) DE.AE-1: A baseline of network operations and expected data flows for users and systems is established and managed
2. Micro-segmentation	PR.AC-5: Network integrity is protected, incorporating network segregation where appropriate
3. Encryption	PR.DS-1: Data-at-rest is protected PR.DS-2: Data-in-transit is protected
4. Multi-factor Authentication	PR.AC: Access to physical and logical assets and associated facilities is limited to authorized users, processes, and devices, and is managed consistent with the assessed risk of unauthorized access to authorized activities and transactions. PR.AC-1: Identities and credentials are issued, managed, verified, revoked, and audited for authorized devices, users, and processes PR.AC-6: Identities are proofed and bound to credentials, and asserted in interactions when appropriate
5. Patching	PR.IP-3: Configuration change control processes are in place PR.IP-7: Protection processes are continuously improved PR.IP-12: A vulnerability management plan is developed and implemented ID.RA-1: Asset vulnerabilities are identified and documented DE.CM-8: Vulnerability scans are performed

NOTE: The following appendices provide information which would be of interest to practitioners and others who would be responsible for implementing a new application-focused approach to security within an organization.

Appendix 2: More Details on Application-focused Capabilities

The following sections give a more detailed and technical explanation of the application-focused capabilities presented earlier in the paper on page 10.

Capability I: Recognize an application and establish a baseline reference for it

This capability enables organizations to better understand their critical applications from a security perspective; determine the components that make up the application, including identifying what services (i.e. software functions) should be running on what servers, and what resources are being used, how should the components interact, etc.

A key aspect is understanding the intended behavior of the application such as:

- What should be running?
- What are the allowable interactions?
- How should the components communicate?

Since applications are dynamic, organizations must be able to track the application as it changes, as developers update the application and as it runs in operation, for example, knowing how many instances are running.

What can organizations do with this capability?

- Protect applications effectively armed with a baseline reference
  - Understand the application and know how to protect it
  - Use one authoritative source of information on the application to configure the application's whole portfolio of security controls
    - Controls assurance and audit teams can also use this reference in evaluating the controls
- Refine privilege so it's tight yet operationally plausible (won't break processes)
  - Have the information required to determine minimum necessary function and interactions for the elements that make up the application–creating a least privilege environment for the application itself.
  - Have the information required to determine minimum necessary communications between system components into, out of, and within individual applications
    - This could dramatically reduce the attack surface.
- Make alerts more actionable
  - Alerts from security tools will identify the application and based on the referenced information, the security team will know how much effort to put into response, how to prioritize, and the remediation options

HOW PRACTICAL IS THIS TO DO?

While in the past it's been difficult to gain visibility into the application and know how it operates, new technologies make it easier.

For existing large applications, technologies are now available which watch the traffic on the network to help understand the components of the application and how they interact.

For newer application architectures, DevOps techniques automate the build process and track all the components used to make up the application from the onset.

ENABLES DEVOPS

Newer applications use DevOps practices and technologies to build, test, and roll-out applications rapidly and frequently.

Organizations can move away from manual security review and testing processes which have long lead times and don't work for agile applications and development processes.

EFFECTIVE SEGMENTATION

The traditional model for network segmentation is based on attributes such as type of server, e.g. web servers or database servers. This can't effectively inhibit lateral movement from one application to another since an application doesn't operate within a single segment of servers, it crosses multiple segments. Attackers can hop across segments. To effectively protect an application requires putting a boundary around it, and having a network control point from which you can control and monitor all traffic going to and from any part of that application.

- Improve signal to noise
  - With tightly-controlled systems, there is far less potential for unauthorized access, function, or interactions; the number of alerts will be fewer
  - Fewer alerts means a much clearer signal: less noise and fewer false alarms
- Move away from continually chasing threats
  - Threats are constantly changing. This approach doesn't rely on having to always understand a new threat in advance.
- Drive closer alliance between security teams and application teams (see sidebar)
- Scale implementation: focus on understanding a small set of critical applications first
  - Start with critical applications that have a relatively small set of components with very specific jobs, for example, "This web server is part of Application X and this process should be the only thing to communicate out to Y."

Capability II. Compartmentalize system components into individual applications

Organizations can use the reference information described above to determine the system components that make up an individual application, then compartmentalize the system components and put a logical boundary around the system components. The boundary provides a way to uniquely define and label a single application. It also confines an attack to a single application, if one application is compromised.

What can organizations do with this capability?

- Set up a single gateway for effectively enforcing policy on the application
- Improve protection of applications from threats within the network
  - Segmentation based on infrastructure, such as type of server is not effective (see sidebar)
- Uniquely identify and label an application (see sidebar)
- Develop one policy to apply to one application
  - By uniquely defining and labeling an application, policy can be applied to it
- Inhibit lateral movement from one application to another
  - If an attacker gets into one application, they will find it difficult to move to another
- Apply application-specific controls at the boundary of the application
  - More critical applications can have higher levels of protection and have greater inspection. Even if a weaker system somewhere in the environment gets compromised, the attacker can't move laterally to the more critical system

UNIQUE IDENTIFIER

Traditional security controls still rely on using the same, static label for the whole stack: from the application to the OS to the hardware. But this doesn't work for modern applications since they don't reside on static servers. VLAN identifiers don't solve for this either; they isolate multiple applications and do not provide a unique identifier for each application. For security controls to be effective, they need to have a unique identifier to use in applying policies to each individual application.

Capability III. Position defenses around each individual application

Using the unique label provided by the boundary of the application, organizations can configure security controls to apply to an individual application. For example, a firewall can be set up to protect a single application using the label of the application's boundary. Organizations can also use the reference information to configure the security controls. For example, using the reference information, the organization can set up an Intrusion Prevention System and develop a rule set for the individual application it is protecting.

What can organizations do with this capability?

- Optimize the placement of policy enforcement points
  - Policies are enforced at the application boundary
- Simplify policies for protecting applications
  - Organizations no longer have to deal with the massively complex policies involved in current approaches which attempt to, for example, use one firewall for protecting thousands of applications (see diagram 4 below)
- Reduce the complexity of managing encryption keys
  - It is much easier to distribute keys for encryption/decryption to the system components of one individual application rather than many applications
- Enforce authentication policies (MFA) using simplified mapping of what applications users and system components' get access to and what factors are required
- Avoid misconfigurations of security controls
  - Controls are set up to protect one application
- Increase the protection for individual applications by adding more controls on the boundary
- Have security controls to work together as a system
  - All controls can be synchronized on one application, using the application boundary to label the application
- Allows controls to be attuned to application dynamics
  - Controls can protect on the boundary as the application moves

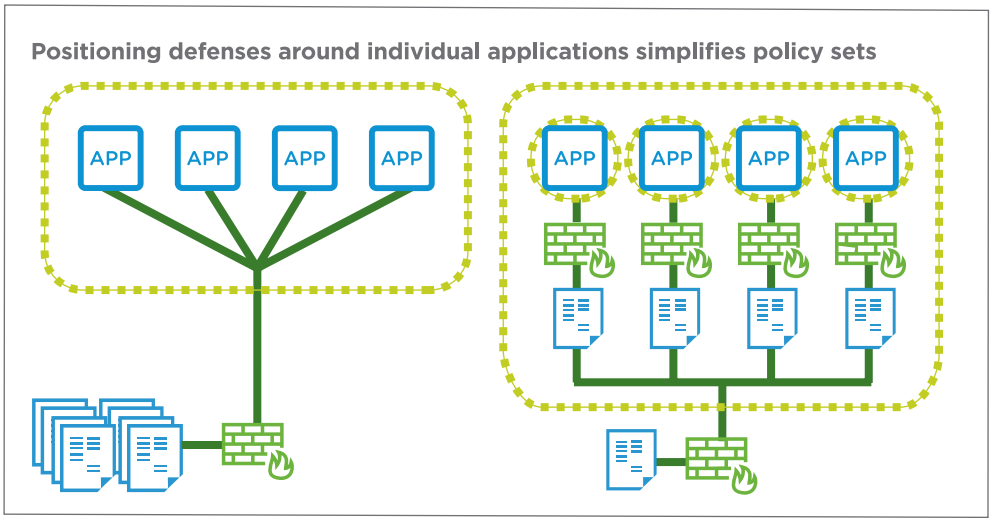

Diagram 4: With current approaches, typically a firewall is set up on the perimeter to enforce policies for traffic to/from all the components of all applications in its boundaries, which may involve 10s of thousands of firewall rules. Policies sets are very large and complex. By setting up a firewall for each application, the firewall need only enforce traffic to/from the components in just one application, vastly reducing and simplifying the policy set.

Appendix 3: Unique Properties of Cloud and Mobile Computing

The application-focused capabilities (described in appendix 2) that organizations can now use to implement a more effective approach to information security have been made possible by recent advances in cloud and mobile computing.

CLOUD COMPUTING UNIQUE PROPERTIES

The basic fabric of the cloud is virtualization, which provides an abstraction layer between the physical infrastructure and the applications.	
Application Context	• The virtualization layer: • Collects, protects and distributes contextual information on all of the applications that are running within the virtualized environment – This is inherent to the function of virtualization, as it controls the dynamics of applications, such as moving workloads to available resources, performing load balancing, scaling up and scaling down of resources as required by the application – It also has a map of all of the workloads and system components within the environment and maintains this map as the workloads are constantly moving around • Has a unique perspective from which to see: – The connection between the application running and the hardware it's running on – The topology of the application • The arrangement of the various system components that make up the application on the network – How the application was provisioned and how it operates at runtime
Isolation	• The virtualization layer: • Provides a separate trust domain – It provides visibility into the guest, but also isolation from the guest • Provides an isolated insertion point to place security controls for protecting the boundary of an application – It maintains the application boundary even if the workloads move to different physical machines or network links
Immutability	• Using the virtualization layer, immutable components can be replaced for every deployment, rather than being updated in-place – A common image can be built once per deployment and can be tested and validated
Software-defined	• With virtualization, the behavior of system components is initialized, controlled, changed, and managed programmatically • It is a flexible control point from which to quarantine machines, reimage machines, block traffic, snapshot machines, insert greater visibility, etc.

MOBILE COMPUTING UNIQUE PROPERTIES

Mobile computing provides unique capabilities via the native functionality of the device, as well as the functionality of virtual desktops and mobile device management technologies.	
User and Device Context	• Mobile computing provides a rich set of data on the user and device to help make risk-based authentication and access control decisions, for example: – Data from users and devices • Biometrics: fingerprint, voice, image • Geographic location • Device ID: serial number, certificate • Network parameters (WiFi, Intranet, etc.) and IP address • Device configuration: hardware, OS, installed applications • Security posture: managed or unmanaged, security software, jailbroken or rooted, status of software updates and patches • Out-of-band: phone call, push notification – Conditional access decisions • Multiple devices: determine if geographic location of smartphone and laptop is different before granting access • Combination of user and device data: determine if it is a trusted user and managed device on a safe network before granting access
Isolation	• Virtual desktops provide an isolated connection to applications – Enables organizations to restrict user access to a particular set of applications (rather than all applications on the network) • Virtual desktops also isolate the use of the applications from the use of the device – Prevents applications and their associated data from being present on the mobile device itself. Instead, the mobile device merely views a remoted display of the applications.
Immutability	• Non-persistent virtual desktops are immutable – They can be instantly created from a controlled master image, then destroyed and recreated with every use. With immutability, it's very difficult for attackers to maintain persistence
Telemetry	• Remote monitoring, policy enforcement and remediation enables: – Continual updates and patching – Erasure of device if it is lost/stolen, fails to check in, or roams outside of secure WiFi – Quarantining or shutting down device if it doesn't meet requirements

Appendix 4: Implementation in the Data Center

This appendix provides specific suggestions for implementing application-focused capabilities in an organization's data center.

Capability	Suggestions for implementation
1. Recognize an application and establish a baseline reference for it	• Create systems of record for how critical applications were setup and the intended interaction among system components – This can be done by working with application teams, looking at provisioning systems, or through learning/baselining or automation systems/blue prints – Serves as critical information of record to spot and diagnose issues • Create a allowlist for an application—for system of components, processes, and how they interact/communicate over a network
2. Compartmentalize the system components into individual applications	• Leverage the virtual fabric to create a logical boundary around the application or service, i.e. micro-segmentation – Enforce that boundary not just with a distributed firewall, but with an isolated L2/L3 network–creating a noncontiguous address space – All the components of the application are contained in a single isolated segment with a single boundary of control • Set up a single egress point – Components of an application inside the segment are free to communicate with each other – Limited set of services communicate across that boundary (DHCP, DNS, AD, etc.) – Application boundary is a defined point on which you can align your controls to inspect the traffic from those servic
3. Position defenses around each individual application	• Use the virtualization layer to align controls to the application – Software-defined networking and software-based security controls now make it operationally feasible to position defenses around each individual application.

Appendix 5: Implementation for End User Computing.

This appendix provides specific suggestions for implementing application-focused capabilities for End User Computing.

Capability	Suggestions for implementation
1. Recognize an application and establish a baseline reference for it	• Use non-persistent virtual desktops instead of persistent applications on endpoint devices as part of ensuring that the applications maintain their intended operation – With a non-persistent desktop image, the operating system and applications are maintained in the intended state by having the desktop image destroyed at logoff and created fresh upon the next logon – In the event that a non-persistent desktop image is compromised, the attack will be destroyed later that day when the user logs off. Attackers typically need time (days or more) to spread the attack from the initial machine through the network so nonpersistent desktops can hinder attackers from moving beyond their initial foothold – It inhibits the attacker from keeping a foothold in your environment. It hinders the ability of Advanced Persistent Threats (APT) to be persistent • Use real-time device security compliance checks to quickly determine if a device is out of compliance with security policies and remediate it instantly or disable its access to corporate resources
2. Compartmentalize the system components into individual applications	• Confine the end-to-end process, i.e. a user connecting to an application – Connect the compartmentalized application to the end user infrastructure • Use virtual desktop infrastructure (VDI) technology to ensure users can only access systems they are supposed to – Using access controls at the application layer – For example, only allow contractors to access the applications they need • When contractors log into virtual desktop, they only get access to one micro-segment (application) • Use VDI technology in conjunction with micro-segmentation to prevent an attacker from spreading an attack throughout the network – The virtual desktop platform can use micro-segmentation to ensure that if a user's machine is compromised (e.g. through spear-phishing), the attacker would only be able to access a small number of hosts rather than thousands • The attacker could only gain access to the limited set of applications that the user can access via VDI – The use of micro-segmentation by the VDI simplifies segregation • It's based on the identity of the user: once user logs on, they are dynamically provided their own view of the network • Don't have to do a lot of complex network mapping ahead of time or preconfigure different sets of desktop pools each having different VLANs associated with them

	• Use mobility technology in conjunction with micro-segmentation to constrain the datacenter resources that a mobile app or mobile device is able to terminate onto – When the device goes to access resources in the datacenter, based on the identity of the device, the device can only access a very limited part of the network, e.g. specific IP or port • Terminate the VPN at the micro-segment boundary providing the user with a secure authenticated connection directly to one application – Traditionally the VPN terminates at the perimeter so that once inside the user has access to lots of places on the network
3. Position defenses around each individual application	• Use VDI technology to apply security controls directly to applications – It is more effective to apply security controls to applications centralized in a data center versus applying controls to 1000s of devices • Leverage OS and application containerization technologies to securely separate corporate apps and data from personal apps and data in a BYOD use case so that corporate security controls can be applied directly to corporate applications – Applications are delivered in a sandboxed and encrypted fashion • Leverage data from endpoint devices to ensure that the level of evidence regarding identity and trust is commensurate with the level of risk of access request to an individual critical application – For example, with an unknown device, the user requires two-factor authentication to access an application. With a trusted, enrolled device, the user proceeds with singlefactor authentication. The device acts as a second factor. Or with a trusted Wi-Fi network (in a corporate office), the user proceeds with single-factor authentication. The network verification acts as a second factor. • Use geolocation information for real-time risk decisions regarding access to a single critical application – Take the geolocation of a user's laptop and smartphone, if they are in geographically disparate locations, the risk level requires additional steps in the authentication process. For example, a push notification could be sent to the smartphone asking for verification. • Use federated identity to make authentication more secure and simplify the logon for the user – Federated authentication to third-party directories avoids insecure practices such as: • Synchronizing your directories • Having multiple passwords which inevitably causes users to write them down or reuse the same password across all applications • Use mobile technology to automatically ensure security readiness of devices regardless of whether devices are running Windows, OSX, iOS, Android, QNX, etc. – For example, poll devices to identify any security issues such as missing patches then remediate by immediately pushing out the patches

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com