Engineering Digital Risk Protections Using Software Defined Perimeter

White Paper

Juanita Koilpillai
Waverley Labs LLC

May 2019

Introduction

The current state of cyber security is slowly transforming from a compliance-based approach to a risk-based approach. Compliance-based approaches gave rise to a plethora of point products that provided specific controls and by design are extremely time consuming and man-power intensive in order to ensure that these products can hang together to provide the necessary protections. Despite all these efforts, hacks are getting more and more spectacular and there is a general sense of despair; so much so that industry experts are going so far as to say that 'cyber security was dead on arrival'.[1] Even with the shift to risk- based approaches, very little has been done to integrate the various point solutions. The net of it is that cyber security can no longer be delegated by the IT security teams alone. It is a business problem and a political issue. This challenge is increasing the need for information security professionals to understand and be able to explain risk from a business perspective. To do this, requires an understanding of the differences between cyber security, cyber risk and digital risk.[2]

The trend is now to elevate cyber security to the executive and the board levels within organizations. This paper addresses digital risk and how the Software Defined Perimeter (SDP) helps to reduce risk from cyber threats. The SDP is a **prescriptive five layer security model** that stops all network-based cyber attacks.

Digital Risk is the risk from the probability of a cyber event and the magnitude of losses they induce. Digital Risk Management provides a **practical, analytical *discipline*** for managing digital risk from a business perspective, by enabling business and technology leaders and their security partners to ***engineer* digital risk** out of business operations and ***mitigate* or *transfer* residual risk**.

The quantification, reduction and control of risk acticvities within an organization will go a long way to help organizations do business online securely. The SDP deployed by Waverley Labs (viz. Panther^{TM} - a commercial version of the open source code developed in cooperation with the Cloud Security Alliance and funded by the Department of Homeland Security) can facilitate reducing risk from cyber attacks and to help organizations 'engineer digital risk' out of business operations. Being proactive helps to simplify meeting FISMA/FedRAMP and other compliance mandates.

[1] https://www.welivesecurity.com/2013/10/15/is-cybersecurity-by-fiat-doa/

[2] see appendix at the end this document

Panther ™

Panther™ collapses the number of individual tools required for compliance controls to obtain the Authority to Operate (ATO). Broadly speaking, the classes of threats that need to be addressed by system owners can be classified into three types of rogue users as shown in the diagram below.

These rogue users make attempts from outside the network perimeter, or they are rogue users who cross boundaries from within the network or they are rogue insiders. By installing Panther™, system owners can ensure that all connections bound for protected servers/services, are made only from authorized users who are associated with validated devices. More importantly, these servers/services are behind a dynamic firewall to essentially be hidden from all network queries and scans from rogue users outside the network and cross domain rogue users inside the network. System owners can then redirect focus on protecting their servers/services from the insider threat. This approach has proven that over 60% of the various FedRAMP

Figure 1

Classification of Threats/Rogue Users

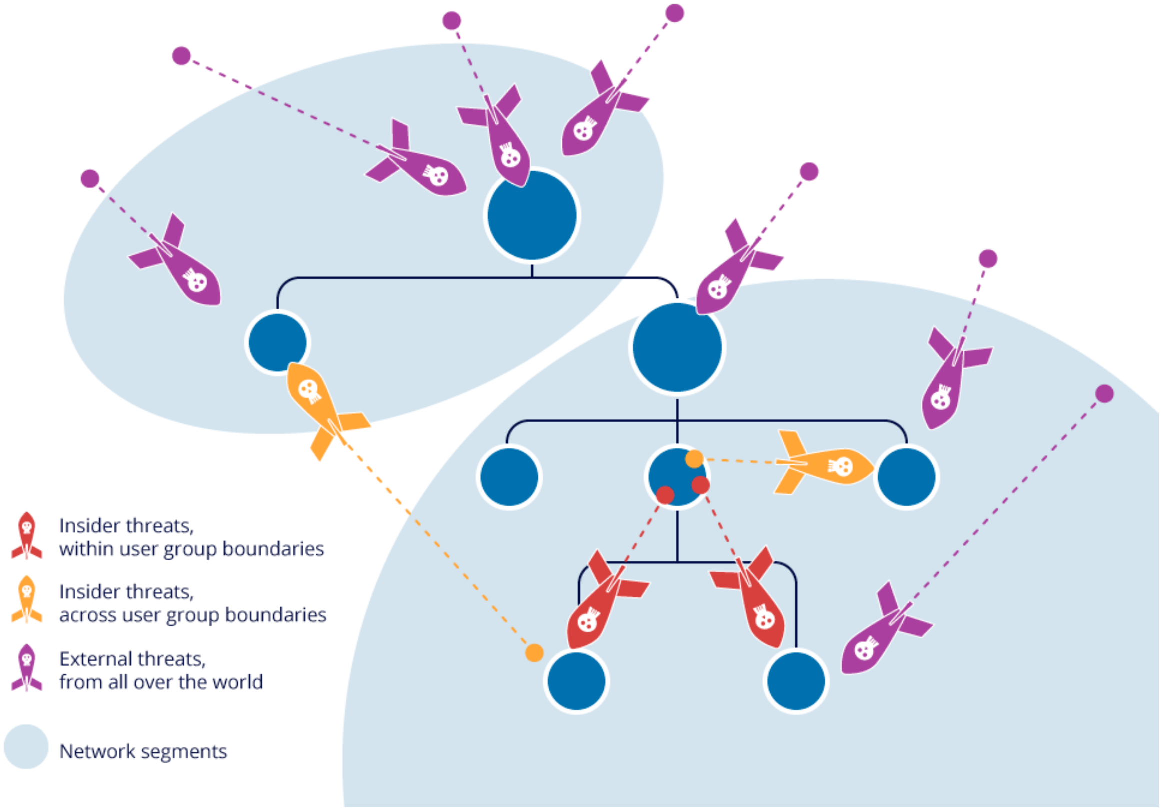

and similar controls will be satisfied and a study by SAIC validates this claim. Moreover, during hackathons hosted by the Cloud Security Alliance (www.cloudsecurityallliance.org), over 10 billion attempts have been made to hack into applications protected by SDP and none were successful.

NIST CyberSecurity Framework Controls Ranking of Panther™ Compared with Other Point Products

	Avg CSF Count	Avg Identify Score	Avg Protect Score	Avg Detect Score	Avg Respond Score	Avg Recover Score	Avg CSF Score
Average Vendor Score	37.23	5.78%	12.97%	13.60%	6.34%	3.67%	12.21%
Panther™	189	71.30%	86.25%	74.58%	72.73%	77.55%	61.97%

FedRAMP & Similar Controls Ranking for Panther™

Control List	Control Count	Total Control Count	Baseline Score
FedRAMP - High	132	407	32.43%
FedRAMP-Low	79	124	63.71%
FedRAMP-Moderate	116	325	35.69%
NIST-CSF	197	263	74.90%

Figure 2

Key Features of a Software Defined Perimeter

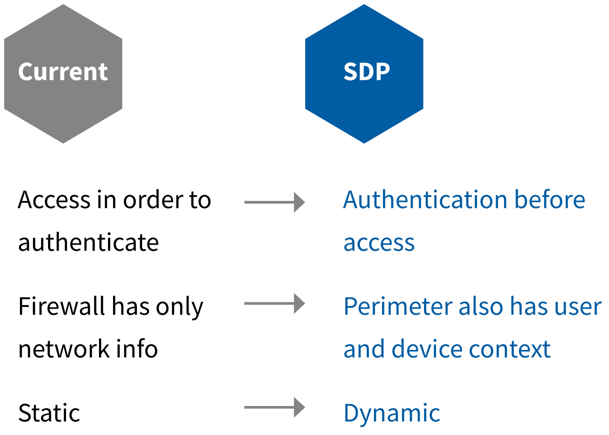

Current	SDP
Access in order to authenticate	Authentication before access
Firewall has only network info	Perimeter also has user and device context
Static	Dynamic

Figure 3

SDP Integration

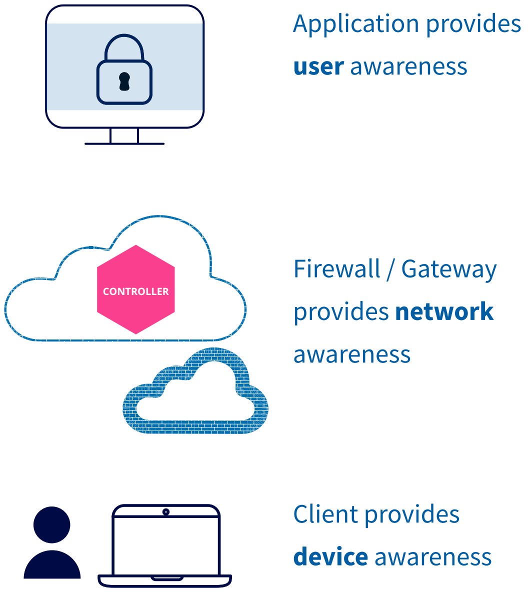

Application provides **user** awareness

Firewall / Gateway provides **network** awareness

Client provides **device** awareness

Today's security products are designed to allow access to servers/services prior to authenticating the users or their devices. This practice requires the enormous responsibility and expense to figure out who the bad actors are and to ensure that systems and servers are securely configured; an equally expensive and difficult task. SDP enforces authorization of users on validated devices prior to connecting securely to servers/services hidden behind a closed dynamic firewall.

Today's firewall is static and ONLY contains information about the network. SDP enables a dynamic firewall that has user and device context to open a firewall for a single secure connection; the firewall remains closed at all other times thereby creating a 'black cloud' instance regardless of whether the connections are made to the cloud or on premise. This feature greatly reduces the complexity of securing applications that organizations want to migrate to the cloud and to comply with FedRAMP and similar controls requirements. The process of ensuring that users are authorized and devices are validated prior to establishing secure connections to hidden servers/services allows the necessary integration of security features and tools that is next to impossible to implement with the way networks and security tools are configured today. By using Panther^{TM}, this process is streamlined to provide an integrated approach to address the security controls required for a FedRAMP and similar audits whether the application or service is on premise or in the cloud or using a hybrid implementation.

Applications typically have the awareness and information about users that the networks do not have. Firewalls and gateways are network aware and have information about the network that applications do not have. Clients have all the information about the devices but none of the users and the network.

SDP uniquely provides the necessary integration of these various pieces of information to become secure. In addition, SDP separates the data plane from the control plane and as a result, this architecture inherently supports all cloud and hybrid architectures.

In summary, as shown in the figure below, the process of securing connections to hidden servers/services includes a token that consists of both the authentication and encryption keys that are verified prior to establishing connections with the protected servers/services. Users are provided tokens via a self-service portal or through existing authentication mechanisms within organizations. Devices are also validated via self- service portals to reduce the management of provisioning.

Figure 4

SDP Overview and Process

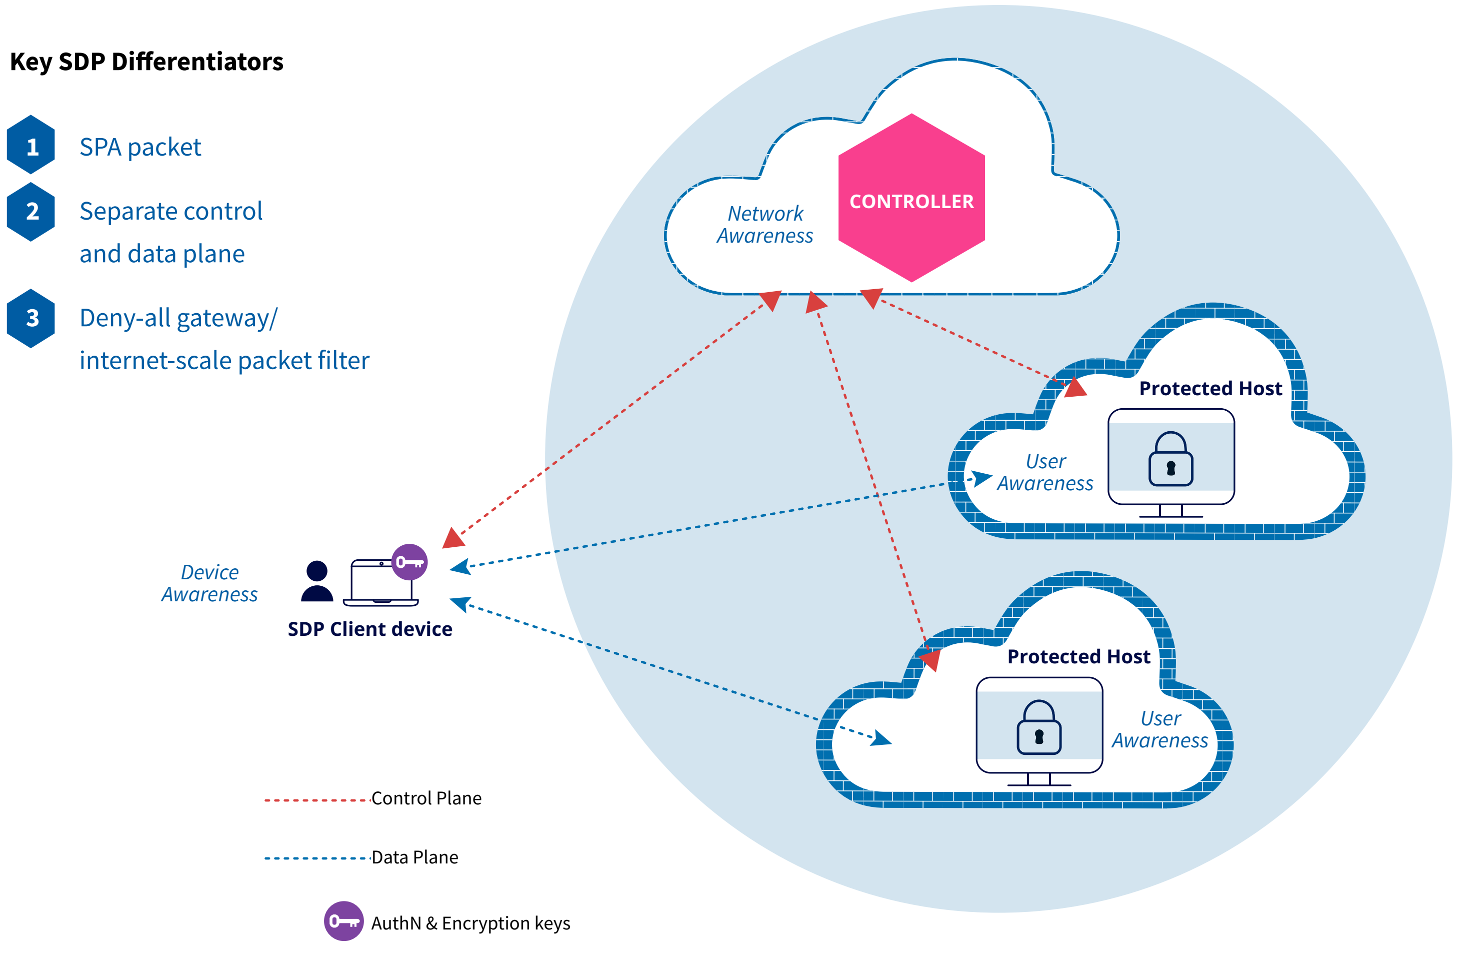

Key SDP Differentiators

1. SPA packet
2. Separate control and data plane
3. Deny-all gateway/ internet-scale packet filter

SDP Process

1. Net facing servers hidden
2. Legit user given unique ID
3. Legit user sends the token
4. Perimeter checks the token
5. Valid device + user = access

In the latest SDP Architecture Document released by the Cloud Security Alliance, several SDP models are expanded upon. Each of these models secure various connections in the enterprise IT stack. Panther™ is designed to address the security of connections in all the models described in the architecture document.

Figure 5

SDP Models Secured

SDP is not IP-Babased
Panther™ orchestrates security
of every connection

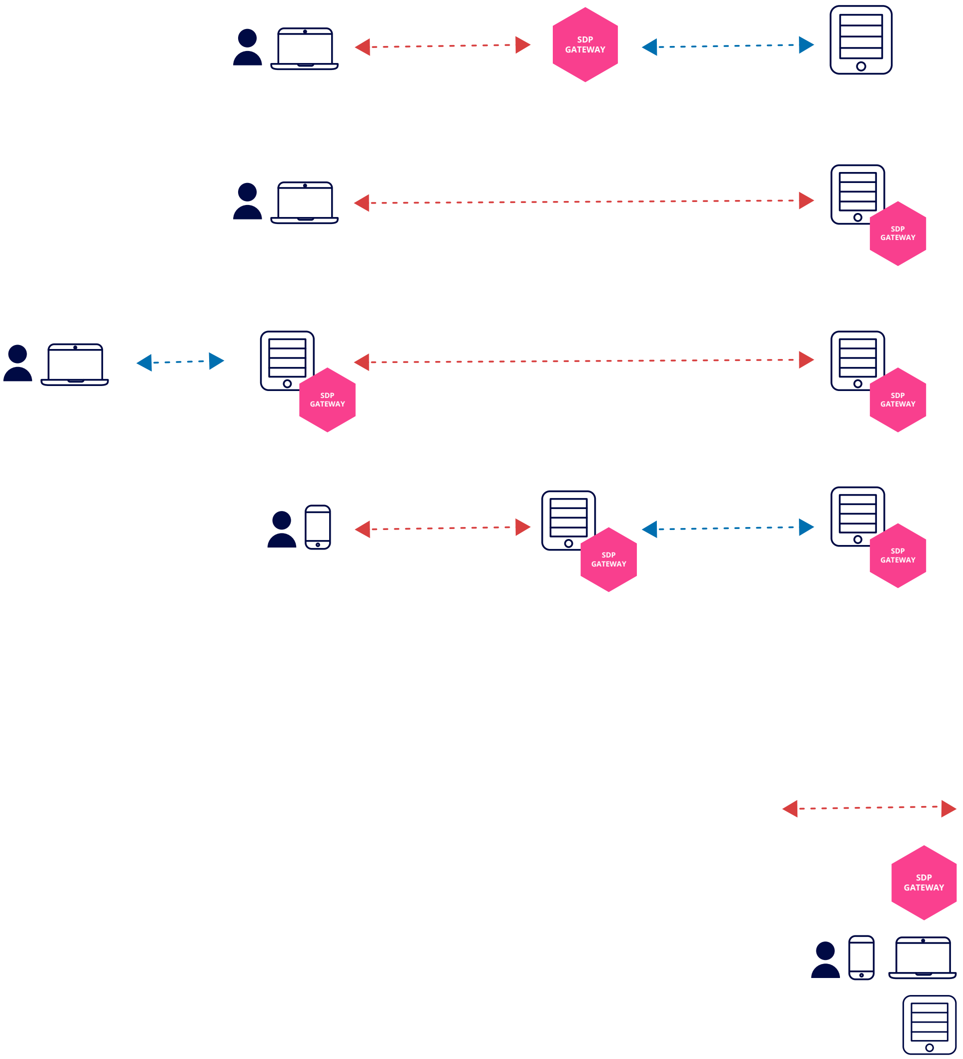

Client-to-Gateway Model
Protect one or more servers

Client-to-Server
End-End to support load balancing, elasticity...

Server-to-Server
API based, SOAP...

Client-Server-to-Client
iPhone, peer-tp-peer...

Secure Connection

SDP Gateway

SDP Client

Hidden Servers

Benefits of Software Defined Perimeter

I. SDP Hides Applications from Unauthorized Users

SDP is a new approach to operationalizing applications in the cloud while mitigating all network-based attacks. It protects both legacy critical infrastructure/IT assets and cloud services of all classification levels. It works by hiding critical infrastructure/IT assets within an undetectable, invisible, black cloud, whether the assets are on-premise or in a public or private cloud, a DMZ, a server in a data center, or even inside an application server. A firewall/gateway combination is placed in front of servers/services that are protected with a deny-all firewall rule-set. Network scans will therefore not show any servers/services present on the network being scanned. Unauthorized users cannot access what they can't see.

II. SDP Enables Applications and Wraps Them in an Invisible Cloud

SDP uses a combination of tried and true security protocols that were previously unconnected until the Department of Defense (DoD) announced them working in concert. The Cloud Security Alliance adapted the generalized DoD workflow but modified SDPs for commercial use and made it compatible with existing enterprise security controls. Where applicable, SDP has followed NIST guidelines on cryptographic protocols (such as Mutual TLS) and securing applications in the cloud. The perimeter around the application includes the users and devices that connect to it thereby easily extending the perimeter to networks in the cloud for each connection.

III. SDP Provides Zero-visibility and Zero-connectivity to all but Authorized Users & Devices

Connectivity is based on a need-to-know access model. Device posture & identity are verified before access to application infrastructure is granted. The application infrastructure is effectively invisible or black. There is no visible DNS information or IP addresses. SDP combines security protocols previously not integrated, they include Single Packet Authentication (SPA), Mutual Transport Layer Security (MTLS), device validation, dynamic firewalls, application binding. The initial SPA packet contains all the relevant information about the users and devices that are verified before every connection. All other packets are dropped as they are by default from unauthorized users. Prior to opening up the firewall for connections, the token in the initial SPA packet is authenticated based upon the authentication level required by the server/application thereby enforcing user authentication policies of the organization for every connection.

IV. How to Enable a Dynamically Provisioned Perimeter

Today, management of firewall rules is an expensive and time consuming process. Some organizations manage as many as 250,000 firewall rules and changes to rules is complicated requiring multiple authorizations in addition to security checks by independent organizations. Placing a software firewall/ gateway combination with a deny all rule-set behind these physical firewalls is the first step to creating a dynamic firewall. Every packet hitting the firewall is inspected for a single packet authentication (SPA), then is quickly verified by the gateway for a connection request. With a single SPA packet, valid requests are processed by the SDP controller, which then creates a rule dynamically on the firewall after the authentication/encryption tokens inside the SPA packet are verified. Once established and the firewall is closed again, connections made are then not seen by rogues outside the network or rogues outside the user domain within the network.

V. SDP Authenticates and Validates Devices in Addition to Users

Network policies are enforced at the firewalls and gateways. However, today it is difficult to enforce these policies at end-points that are mobile without additional tools (such as mobile device managers), while at the same time securing the perimeter. While VPNs and VLANs can bring these devices into the network securely, connections from the clients to applications within the network cannot be easily secured because the networks don’t have that end-point visibility. With SDP, devices are validated by using embedded clients and those devices that do not comply with the desired network policy will not be given access to the network or have reduced visibility to points within the network.

VI. SDP Works with User Authentication (PIV card, ID management) Systems

In addition to network policies, FedRAMP and other guidlelines require user and device policies that are hard to enforce by stitching together the myriad of security tool options in a streamlined and efficient way. Waverley Labs has been working on the NIST Federated Identity Management Guideline for cloud security and have defined workflows to ensure that the credential management and associated policies for both the user and user environment (i.e. device) can be enforced. SDP’s Controller is the central authority that is scalable in its efforts to validate users by integrating external sources of information such as Active Directory (roles, permissions etc.), PIV Card Readers (keys, checksums) and the like to ensure that the organization’s user security policies are enforced prior to allowing encrypted connections to protected servers/services.

VII. Types of Attacks SDP Prevents

Waverley's SDP solution focuses on securing connections and strengthening the underlying structure rather than adding to the instability of the systems. By using concepts described above, SDP protects mission critical services and applications from the top OWASP attacks (https://www.owasp.org/index.php/Category:OWASP_Top_Ten_Project)

In the diagram below that represents a typical cloud application, SDP is shown to not only stop all network-based attacks on protected servers, but it also handles credential theft, man-in-the-middle attacks and vulnerable code, and server exploitations equally well with minimum disruptions.

Figure 6

Client/Server Apps or Cloud Apps, Attack Vectors

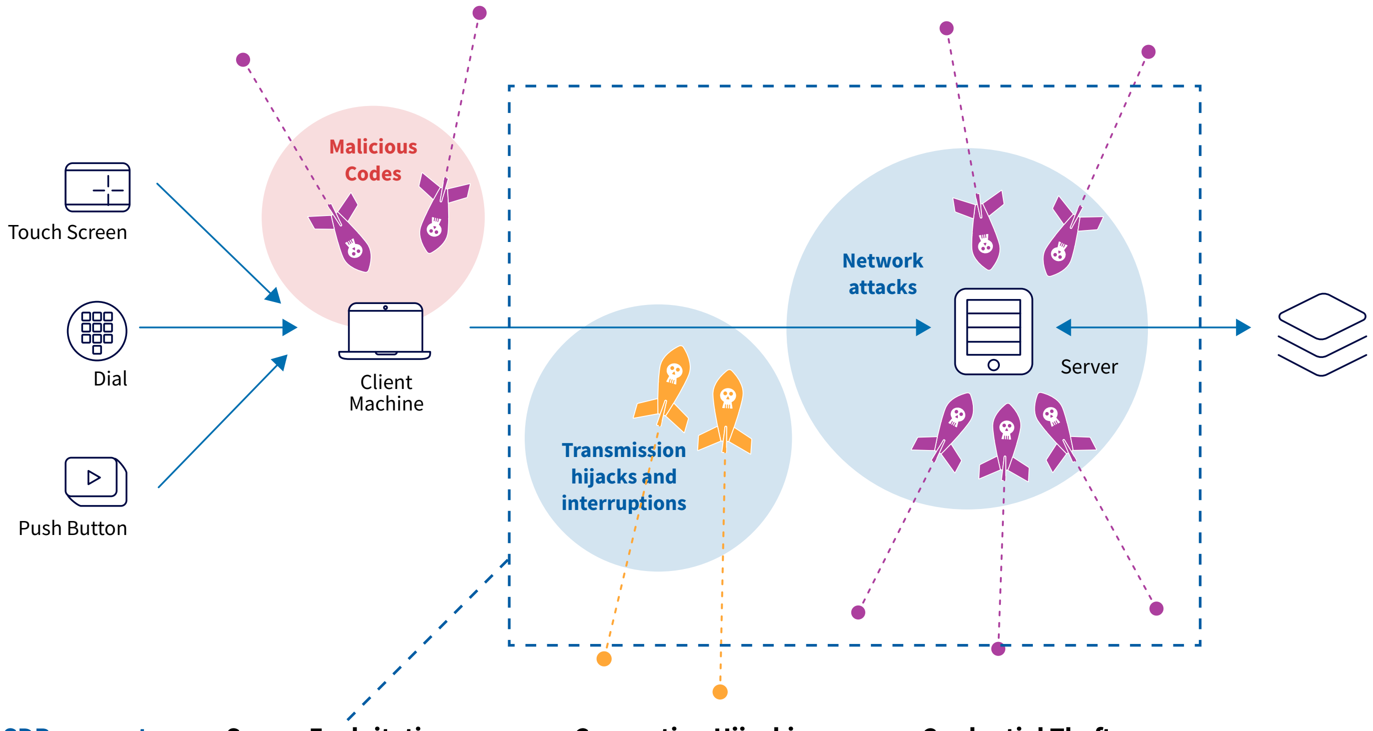

SDP prevents these attacks	Server Exploitation	Connection Hijacking	Credential Theft
	Misconfiguration Vulnerabilities Injections Denial of Service	Man-in-the-middle Certificate forgery DNS poisoning	Phishing Keyloggers Brute force
With these features	Server Isolation SPA Dynamic Firewall	Encryption Pinned certificates No DNS	Multifactor Authentification mTLS Device Fingerprinting

VIII. SDP Orchestration Provides Incident Response in Seconds

Because of the various security technologies that SDP tightly integrates, system owners can now have the ability to control how secure their applications/services are. They can disconnect users in seconds, block connections in seconds, shut-down devices in seconds. All this is possible because any packet hitting the SDP without a valid SPA packet is deemed rogue. Incident response is now focussed on not just analyzing data that SDP provides but they have an added advantage of knowing all the good connections so that they can deal with the rogue ones in seconds. SDP uniquely provides this ability to system owners based upon predefined policies set by the organization as a whole.

XI. SDP Reduces Costs

SDP meets the FedRAMP requirements with a much less complex network security approach. VPNs can now be discontinued for a better and more flexible approach. Two-factor authentication can now be more easily implemented and integrated with the network and applications. The security team no longer needs to manage large lists of rules on the firewall or review, analyze and classify large volumes of log data. SDP makes data collection for forensics easier. SDP provides simpler IT Security, vulnerability scans on networks no longer required as frequently and applications can be hardened in the cycle. This approach ensures network segmentation by default, while handling the OWASP TOP 10 threats and makes logging and auditing simpler.

Appendix

The selection of controls to become compliant to NIST RMF places onerous requirements on business/system owners to implement. While most (typically 75%) of controls can be implemented organization-wide, the implementations require multiple security products that are hard to maintain let alone integrate into a holistic security architecture. Similarly, creating a profile for the NIST CSF using the list of controls gets complicated and most business/system owners don't have the resources to use these profiles to secure their businesses. Instead, business/system owners should focus on day-to-day operations and engineer digital risk protections at all layers of the network - which they can do using SDP. By defining failure and/or risk scenarios (for example, protecting gauges on gaslines from remote cyber attacks), quantification or losses occurring from these risks help appropriate the budget to keep them the worst scenarios from occurring. What controls to prioritize and implement and analyzing gaps to close becomes a more productive exercise. Providing the much needed artifacts to assess the controls and measuring progress while responding in a timely manner to potential attacks is the cornerstone of digital risk management and SDP goes a long way accomplish this.

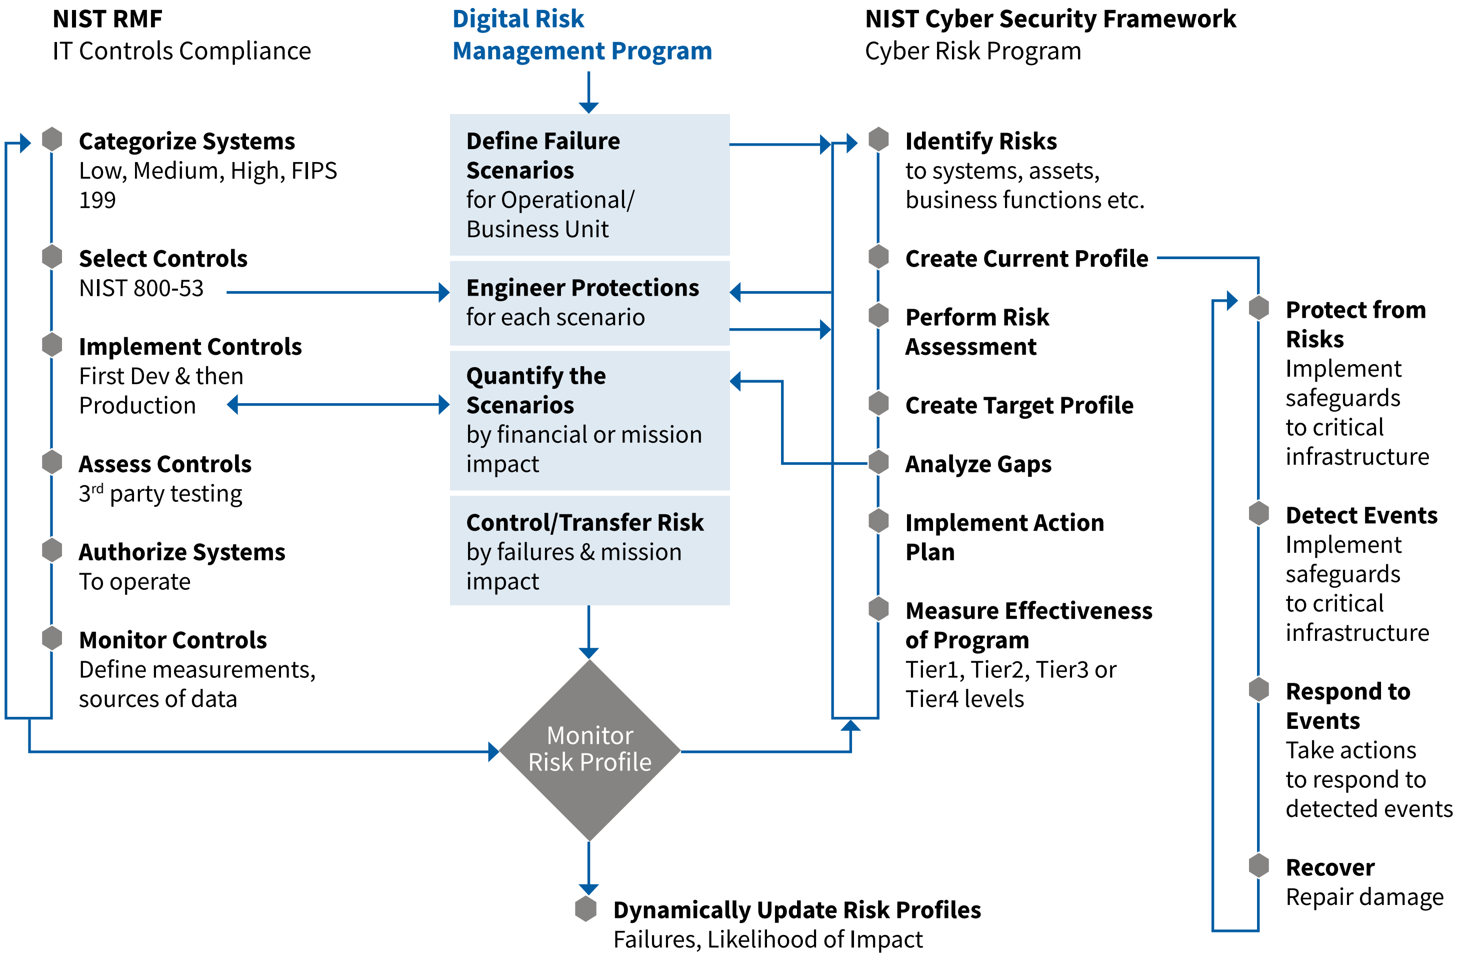

NIST RMF
IT Controls Compliance
Digital Risk Management Program
NIST Cyber Security Framework
Cyber Risk Program
Categorize Systems
Low, Medium, High, FIPS 199
Select Controls
NIST 800-53
Implement Controls
First Dev & then Production
Assess Controls
3rd party testing
Authorize Systems
To operate
Monitor Controls
Define measurements, sources of data
Define Failure Scenarios
for Operational/ Business Unit
Engineer Protections
for each scenario
Quantify the Scenarios
by financial or mission impact
Control/Transfer Risk
by failures & mission impact
Monitor Risk Profile
Dynamically Update Risk Profiles
Failures, Likelihood of Impact
Identify Risks
to systems, assets, business functions etc.
Create Current Profile
Perform Risk Assessment
Create Target Profile
Analyze Gaps
Implement Action Plan
Measure Effectiveness of Program
Tier1, Tier2, Tier3 or Tier4 levels
Protect from Risks
Implement safeguards to critical infrastructure
Detect Events
Implement safeguards to critical infrastructure
Respond to Events
Take actions to respond to detected events
Recover
Repair damage

About the Author

Juanita Koilpillai is Founder and CEO of Waverley Labs a cyber risk engineering company. She has spent 30 years developing systems in computer security, network management and distributed software. She is currently pioneering the field of digital risk management and is the technical advisor to the Digital Risk Management Institute. As part of that effort, she is leading the open source software-defined perimeter (SDP) effort for “black” apps in the cloud with the Cloud Security Alliance. The SDP is a prescriptive five layer security model that stops all network-based cyber attacks.

She is an active contributor to NIST and led the creation of a security risk index system for moving apps to the cloud (NIST 500-299). She was a key member of FEMA’s Enterprise Security Management Team and has served as Principle Investigator for several Department of Defense initiatives. She co-founded CyberWolf, one of the most advanced automated attack sensing and warning systems that was deployed by government and later acquired by Symantec.

About Waverley Labs

Waverley Labs is a leading a leading provider of digital risk management software and services that helps large organizations reduce their exposure to digital risk. Its products and services range from the industry’s first open source software defined perimeters (SDPs) for large federal agencies, to the assessment, quantification, and mitigation of digital risk from the business perspective. Waverley Labs’ automated analysis and visualization capabilities provide business leaders, risk officers and CISOs with an at-a-glance view of business risks prioritized according to business impact and recommended risk mitigations. Waverley Labs works closely with NIST and the Cloud Security Alliance to provide thought leadership in digital risk management. For more information visit http://www.waverleylabs.com.