



# MANAGED DIGITAL SECURITY WHITEPAPER

THE BLUEPRINT FOR BEST PRACTICE DIGITAL  
PROTECTION CAPABILITY WITH AWS AND f5

Author: Simon Morse, Security Practice Director



# TABLE OF CONTENTS

Introduction	2
Our architectural framework	6
Edge security	8
How to leverage preventative capabilities at the edge	10
Fine-grained filtering	11
Security best practice checklist	12
Application onboarding	15
Security best practice checklist	16
Executing tests are intrinsic to pattern development	18
Conclusion: Defence-in-depth for digital security efficacy	20
About Versent	22

# A PROVEN, REFERENCE ARCHITECTURE FOR BEST PRACTICE DIGITAL SECURITY



**Forbes estimates that 83% of enterprise workloads are already in the cloud. As digital is now the norm, the most common threat vectors are web-based: zero-day exploits; brute force attacks; trojans; phishing; ransomware; DDoS; compromised credentials.**

The purpose of this whitepaper is to provide a best practice reference architecture for the safe and efficient adoption of digital protection

capability using components deployed in AWS and augmented with specialist security capability from f5. The recommendations and considerations outlined have been drawn from our various customers' requirements, and then battle-hardened and refined by Versent in a pre-fabricated solution that rolls out the controls quickly and allows efficient ongoing maintenance via automation.

Our architectural framework will help enterprises to defend against a range of attack techniques against internet accessible applications. The technology examples are described using AWS components, but the architecture is designed to be cloud agnostic. At a minimum it needs to allow downstream connectivity to a heterogeneous set of Origin Servers – either deployed in AWS, or alternative cloud platforms (such as Google Cloud Platform and Azure), as well as on-premise environments.

## The focus of this blueprint is to –

- Address real world concerns that organisations will need to defend against. We breakdown actors and threats, and outline a **layered approach to security controls**.
- Outline the implementations of controls and capabilities that are used to **defend against external threat actors**<sup>1</sup>.

---

<sup>1</sup> *That is, threat actors are typically acting with low or zero levels of both privilege (at an application or infrastructure level) and application knowledge or context, then and seek to gain a foothold in the organisation or directly dupe users of the system into revealing information.*

# DIGITAL SECURITY BEST PRACTICES & CONSIDERATIONS

## OUR ARCHITECTURAL FRAMEWORK

The reference architecture describes the recommended layered approach to handling malicious traffic, for higher efficacy and economy. It seeks to apply cheaper and less targeted defences, to eliminate crude attack types early in the network path, and then apply more targeted and computationally expensive defences to counter more sophisticated attack techniques further down the network path.

The philosophy is to progressively optimise the «signal-to-noise ratio» towards the network core and allow per-application policy to be implemented in an efficient and cost-effective manner. It is also intended that the defences work cohesively to defend against threat techniques (the layered approach), and that a given detection / defence technique may be effective against multiple threats.

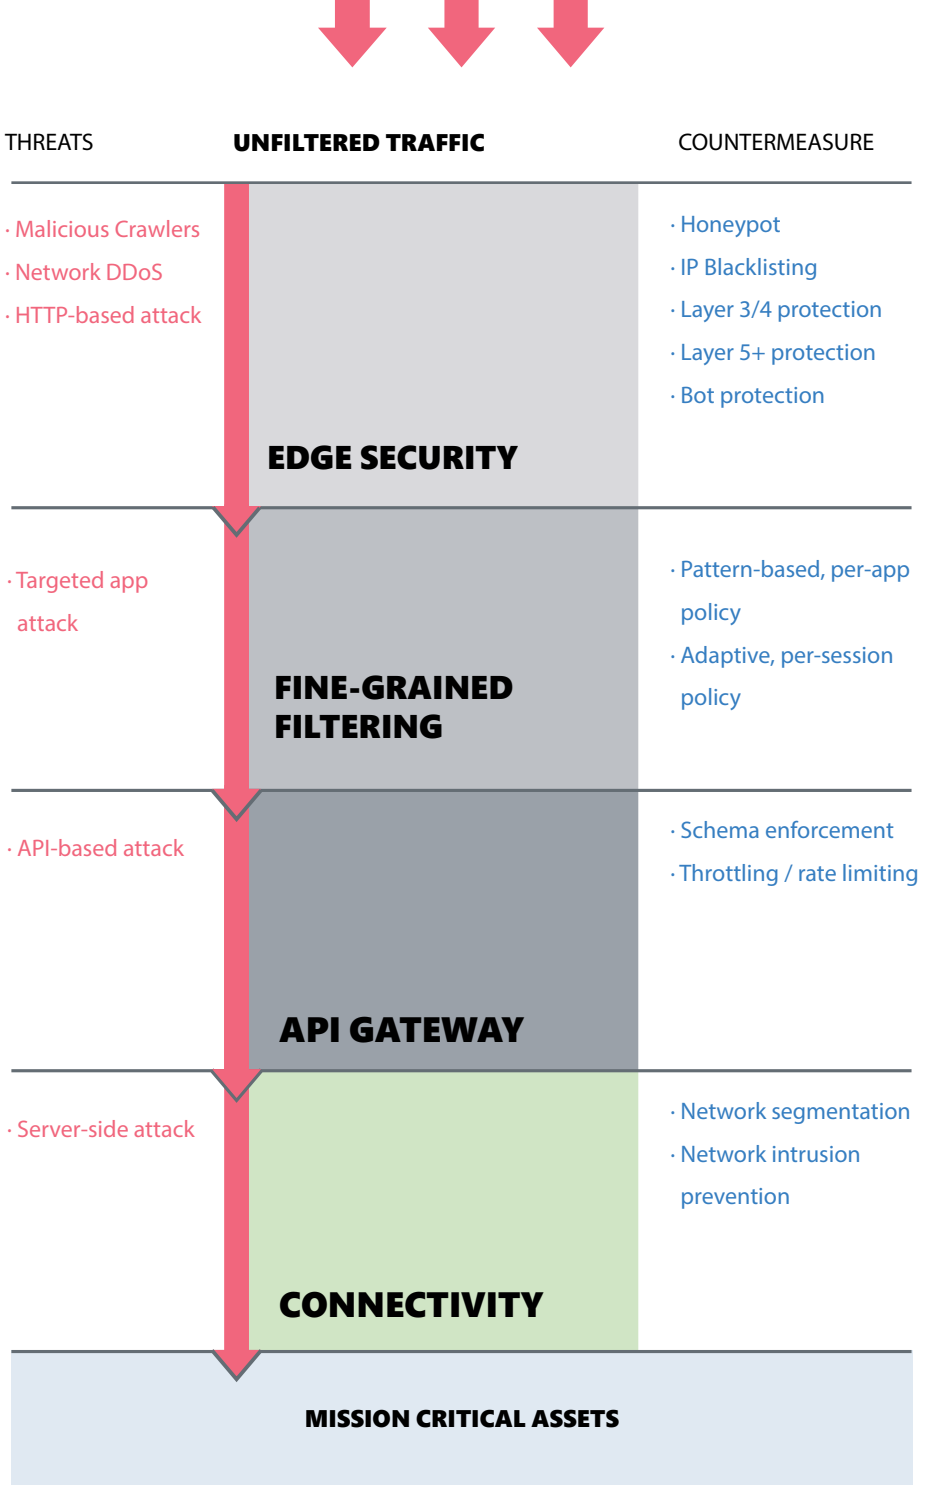


Figure 1: Conceptual model of how the various preventative capabilities work holistically to provide the digital security solution.

## Key features of this approach<sup>2</sup> :

- The filtering capability at each stage is not overambitious - rather that they work to progressively defend against threats
- A variety of deployment topologies, best practice DevOps pattern-based onboarding and environment management are fully accommodated for
- Focus is on the up-front defensive measures, but the architecture is modular (the various API and Integration-based countermeasures are included for context)

## EDGE SECURITY

As Edge capabilities evolve and have richer features, there is a necessity to provide protection further out from the application core. It is also desirable to filter malicious traffic as early in the attack chain as possible. This allows more targeted defensive techniques to be applied after clearly erroneous traffic has been eliminated.

---

<sup>2</sup> There are preventative techniques that are used to defend against direct attacks on the enabling infrastructure itself (security groups, patching, etc.) For simplicity, these have been omitted, as they do not directly defend against the data or downstream organisational assets.



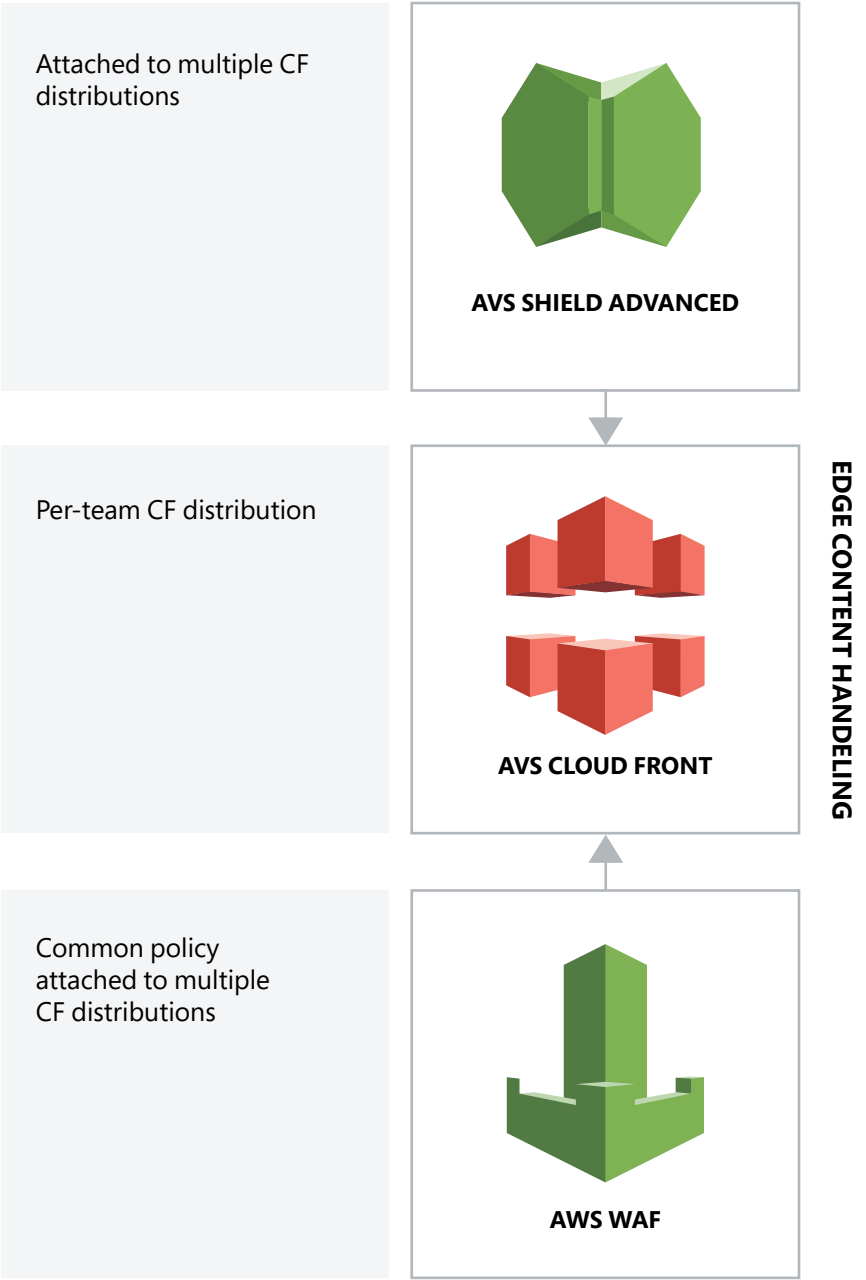





Figure 2: Bundling AWS products for enhanced Edge security

# HOW TO LEVERAGE PREVENTATIVE CAPABILITIES AT THE EDGE

 **IP Blacklisting** can be employed to immediately eliminate traffic from known bad source IP addresses. This can be seeded with intelligence feeds from Amazon or other security vendors. Blacklists can be implemented as a simple AWS WAF rule with a lambda to regularly refresh the list.

 **DDoS Protection** is available at low cost from AWS via their Shield service. The paid Shield Advanced service leverages the AWS DDoS Response Team. They proactively take action against a wide variety of DNS and Layer 3/4 attack techniques such as volumetric and Low-Slow attacks.

 **Edge WAF** can be employed to filter common HTTP based attacks at the edge. Care should be taken not to be too aggressive in this policy, but rather to have a broad and consistent policy that is focused on patterns of bad traffic rather than attempt to catch targeted application attacks.



**Honeypots** can be deployed to try and divert attacks. They can be implemented at an endpoint level where user agents that connect to a honeypot can be used to dynamically update the IP Blacklist. A honeypot can also operate via decoy HTTP headers that have no application utility, so that when attempts to manipulate them are detected, they are clear indications of a malicious user agent.



**Bot protection**<sup>3</sup> is best presented at this layer so that this traffic is eliminated *before* invoking fine-grained capability or bypassed completely for low risk origins such as those with limited input handling. This topology works well with solutions such as Shape from f5, although for some solutions where the technology in question is more tightly coupled to the fine-grained filtering, this can result in routing traffic unnecessarily through the heavy-weight filtering capability. Solutions typically utilise a combination of client-side JavaScript injection and / or integration with external SaaS-based threat intelligence services. Mobile / API based solutions will need to integrate client-side libraries to perform similar functions.

---

<sup>3</sup> It should be noted that Bot detection is primarily focused on classification of user agent behaviour as opposed to attempting to defend the application in any way.

# FINE-GRAINED FILTERING

The fine-grained filtering component is the heaviest component of the digital security architecture, as the intention is to safely remove as much traffic upstream in the Edge Filtering component as possible, before invoking specialist capability of this type (which is typically computationally more intensive, thereby attracting higher licensing charges).

## SECURITY BEST PRACTICE CHECKLIST

1. Apply **pattern-based enforcement** of the WAF policy, as well as site-specific policy that is able to be tuned independently (as per the Application Onboarding below)
2. Have a disciplined set of **engineering pipelines to handle upgrades** of the core technology stack and development / maintenance of new patterns for adoption by tenant applications

3. Develop a **toolkit for the operational maintenance** of the service including regression suites to exercise all the supported application patterns, including interaction with the API gateways where relevant
4. Fine-grained filtering should be performed *before* traffic is forwarded to an API gateway
5. Augment your core fine-grained filtering capability by **obtaining insights from SaaS endpoints**, e.g. for consumption of threat intelligence feeds, statistical services, or dynamically updated rules that are able to defend against emergent threats<sup>4</sup>

## KEY CONSIDERATIONS

1. The implementation approach for this layer needs to be adjusted based on the **deployment approach** taken by the organisation – common approaches are to centralise the deployment of this capability, as many solutions are licensed based on deployed footprint rather than throughput. Alternate approaches are to deploy the

---

<sup>4</sup> Techniques such as this are sometimes referred to as virtual patching

protection capability close to the origin servers (which allows for operational independence); or a hybrid approach. The AppProtect module for NGINX+ delivers the f5 WAF capability on a flexible runtime deployment platform and allows for any of these deployment approaches to be supported with automation of infrastructure orchestration and policy management.

2. The **level of operational maturity** within application teams will affect the degree to which the architecture and operational model must be centralised or is able to be safely decentralised and delegated to DevOps teams to operate.

# APPLICATION ONBOARDING

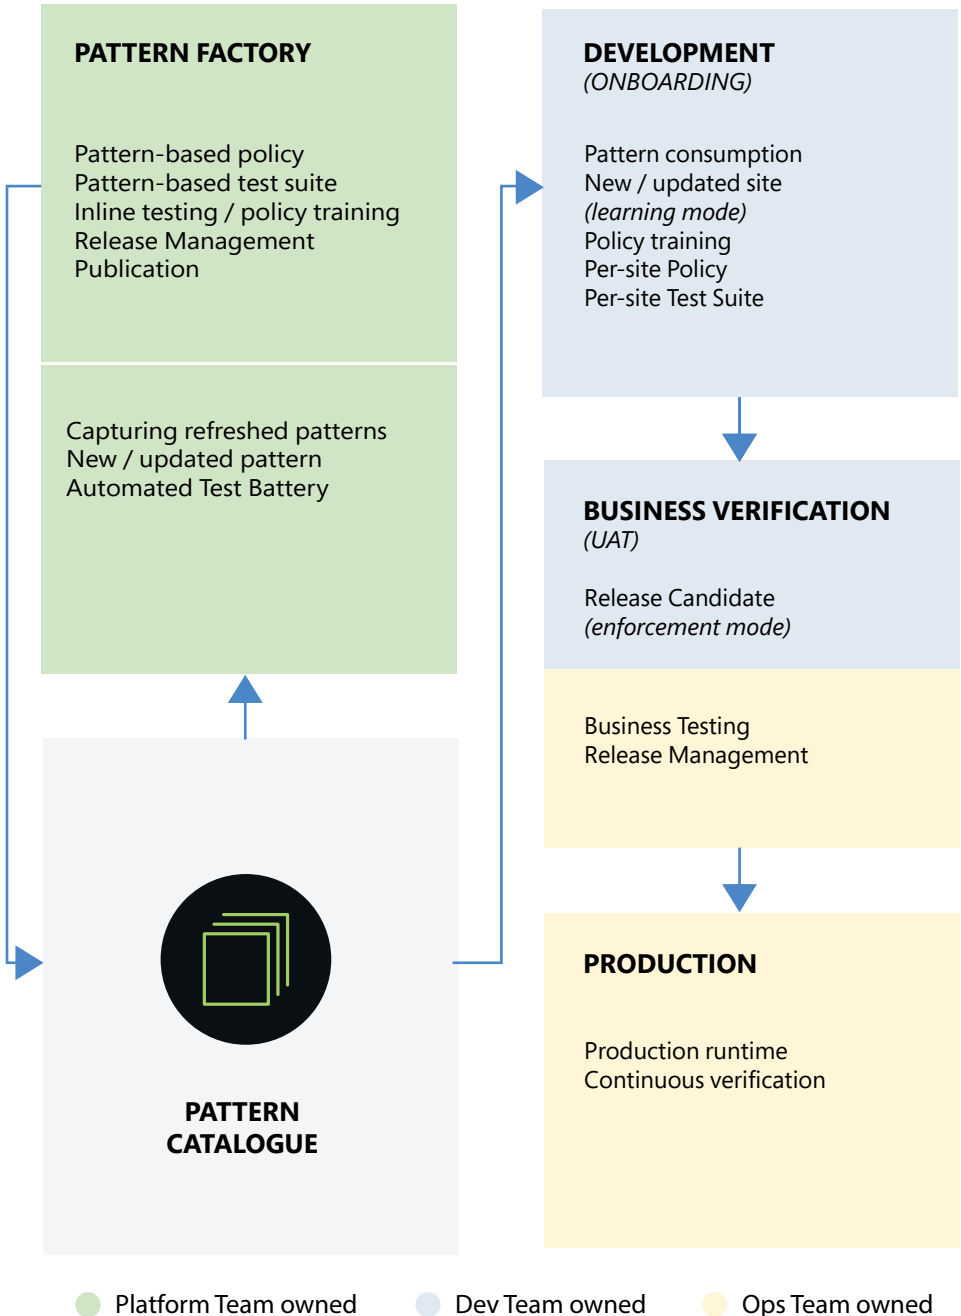


Figure 3: Centralisation of pattern development & use of pre-defined patterns expedite application onboarding

# SECURITY BEST PRACTICE CHECKLIST

- 1. Onboard applications using pre-defined patterns** (maintained by a central team, Security or Cloud Centre of Excellence team, or similar). The patterns will need to be geared to common web technology frameworks (e.g. Wordpress, Joomla, Drupal, .NET etc.) that have common approaches to URL pathing, parameter / header usage, behaviour, navigation and session management. The established, common approaches to identifying normative interactions with the site can be leveraged as the basis for defining fine-grained filtering policies.
- 2. Identify group / function responsibility** for the different phases of the operational process. In the example above, a centralised platform team publishes a number of patterns, individual tenant teams adopt and adapt these, and a centralised ops team manages the production environment.



3. In addition to the Production environment running the live application, a **separate pattern factory environment** is required where the patterns can be developed and maintained.
4. For the adopting development teams, an **onboarding environment** is required to tune the pattern for the particular site being maintained.
5. A **formal testing environment** where UAT as well as other non-functional testing can take place, is recommended to isolate these activities.

## EXECUTING TESTS ARE INTRINSIC TO PATTERN DEVELOPMENT

A key aspect of the approach is that the policies can be tested in both a positive and negative sense:

- **Positive / active test cases** - do attempts to interact with the site in an unexpected manner (including well known attack techniques defined by [MITRE](#) and [OWASP](#)) result in traffic being actively blocked and not passed on to the origin servers?
- **Negative / passive test cases** - does a battery of «typical» site interactions (potentially against a “dummy” site) execute without resulting in false positives?

Both sets of tests need to be executed as part of the pattern development, but also to reconfirm the suitability of both the test suite and associated policy in the event of engineering changes, such as:

- Changes to the application framework
- New versions of the WAF engine

- Kernel upgrades in the underlying execution environment that may result in changes to the way in which HTTP connections are handled, or
- New filtering techniques / approaches as the threat and technology landscape evolves

The test cases are published as key elements of the pattern to allow adopting teams to replicate and tune their tests. The goal of the pattern factory is to act as a centre of excellence in the production of best practice, test driven policy development, and to support teams adopting the pattern.<sup>5</sup>

---

<sup>5</sup> A key design principle is that the baseline policy is not too restrictive and result in undesirable rates of false positives. Metrics should be tracked on the false positive rate to ensure that this is the case.

# CONCLUSION: DEFENCE- IN-DEPTH FOR DIGITAL SECURITY EFFICACY

**In our blueprint for security best practices for AWS environments, the tenet of defence-in-depth, or layered security, is central.**

An array of digital countermeasures is leveraged to defend against malicious external activity throughout all stages of the deployment cycle.

To economise, coarse-grained, upstream defences are deployed as a first-line defence, with targeted countermeasures in place to block more sophisticated threats downstream in the network path.

Data breaches cost enterprises an average of **USD3.92 million**<sup>6</sup> - it is prudent to cast a wider net, and have multiple lines of defences, against digital-based attacks.

---

<sup>6</sup> Source: [CSO Online](#)

## **DIGITAL SECURITY – ALL IN-THE-CLOUD, FULLY MANAGED FOR YOU**

Versent's unified managed WAF and CDN digital security solutions are underpinned by AWS infrastructure and f5, with full on-site engineering and support provided. Our Managed Digital Security solutions are designed to be plug and play via CI/CD pipelines.

## ABOUT VERSENT

Versent is an Australian-born technology company, focused on architecting, building and operating cloud native applications, data streams, platforms and services. Our solutions are centred around AWS and best of breed technology. With diversified and deep expertise in professional services, managed services and product, Versent has expanded regionally to Singapore. We have been AWS' ANZ Partner of the Year for three years running (2017, 2018, 2019).

Versent is the only APAC Premier Partner dedicated to AWS, and holds the following AWS consulting competencies: DevOps, Migration, Managed Services Provider, Public Sector, Security, Solution Provider and Well-Architected.

**Contact us to find out how you can effectively block web-based attacks:**

[info@versent.com.au](mailto:info@versent.com.au)

[versent.com.au](https://versent.com.au)

+61 3 8374 7662