Pitwall

Messaging Guidebook

Internal Use Only | Not for External Distribution | Version 1.0 | 2026

> *This document is the single source of truth for how Pitwall communicates. Every piece of external content, every sales conversation, every partner brief, and every social post should be rooted in what is written here. This is not a style guide. It is the story we tell, the rules we follow, and the standards we hold ourselves to. It does not get shared outside the company. A cutdown version exists for external agency partners.*

SECTION 1: BRAND PERSONALITY

Pitwall is a company of pragmatic operators. We were built by people who ran SOCs, managed stacks, worked investigations, and lived the problem we now solve. That background shapes everything about how we communicate.

We do not talk like a vendor. We talk like a peer. We do not impress with polished language or bold claims. We earn trust by being specific, honest, and direct. We speak to security professionals the way one experienced operator speaks to another.

We are confident without being arrogant. We have built something real and we know it. But we do not need to tell everyone how good we are. The outcomes we deliver do that work for us.

We never position ourselves as the answer to a problem the buyer created. The alert flood is not their fault. The stack grew because the threat landscape grew. They made the best decisions they could with what was available. We are here to help them get more out of what they have built, not to second-guess the decisions that got them here.

> *One-line personality test: If a sentence sounds like it came from a vendor brochure, rewrite it. If it sounds like something a trusted colleague would say across a conference table, keep it.*

SECTION 2: THE NARRATIVE

The narrative is the full story. It is not a tagline, not a positioning statement, and not a list of features. It is the answer to four questions that every piece of Pitwall content must be rooted in. Read this before creating anything.

The Problem

Security stacks keep growing. Every year, new threat vectors emerge and new specialized products are built to address them. Each one adds value in isolation. But over time, the cumulative weight of 25, 50, even 100

deployed controls creates a problem nobody planned for: a daily flood of alerts that buries the teams responsible for finding real threats.

The good news is that a comprehensive security stack means security teams have more visibility into their environment than ever before. The unintended consequence is that with so many controls generating signal, many of them misconfigured or generating duplicate alerts, SOC teams spend the majority of their time chasing false positives instead of working real investigations.

The market's response has been automation. AI SOC agents are being deployed to handle triage, and they help. But they are treating the symptom. The underlying issue, that many of the controls generating those alerts are not tuned to the environment they are protecting, goes unaddressed. You can automate triage all day. If the alerts being triaged are not worth triaging, you are automating chaos.

SOC teams want their stack running like a finely tuned F1 car. The problem is they are fighting fires every day and never have time to get under the hood. There is no easy way to see, in real time, how every control in the stack is performing. There is no automated way to diagnose what is causing a control to underperform. So the alerts keep coming, the analysts keep chasing them, and the real threats keep having more room to hide.

The Way Out

Pitwall gives SOC teams the one thing they have never had: real-time visibility into how every control in their security stack is actually performing, and the ability to fix what is not working without adding headcount or touching their architecture.

Pitwall automatically logs every alert generated across the entire stack, from the moment it is generated to its ultimate resolution. Over approximately the first week, Pitwall begins identifying trends, showing which controls are operating outside accepted standards. From the dashboard, a SOC manager can click into any underperforming control and ask Pitwall to auto-diagnose the issue. Pitwall generates human-readable output that explains exactly what is happening: why a specific control is generating noise, what configuration change would fix it, and whether Pitwall should make that change automatically or flag it for human review.

As the team interacts with Pitwall over time, Pitwall learns their environment and their preferences. Recommendations get sharper. The feedback loop gets tighter. The stack gets better, continuously, without anyone having to carve out time for a dedicated audit.

Early adopters are seeing false positive reductions of up to 45% within 30 days. That is not a product claim. It is what happens when a security team finally has the visibility to tune what they have already built.

The Credibility

Pitwall was built by Joe Stack and Bill Argentina, two people who spent their careers inside the problem we now solve. Joe spent over a decade building security stacks in finance and healthcare, starting as a SOC analyst before moving into architecture. He has lived through SIEM correlation rules that brought an entire SOC to its knees. Bill came up through the business side of security, serving as a CISO and learning firsthand what it feels like to sit in front of a board and try to justify a security investment with no data to back it up.

They did not build Pitwall because they studied the alert fatigue problem. They built it because they could not find a solution when they needed one. That is the only reason this company exists.

The Political Safety Valve

No security team needs to rip out their stack to use Pitwall. No infrastructure changes are required. No dedicated internal resources are needed to run it. The entry point is the Pitwall Shakedown: a 14-day, no-obligation assessment of up to 20 security controls that runs against a live environment and delivers a boardroom-ready report at the end.

The Shakedown was designed specifically to remove the political risk from evaluation. A SOC manager can start it today without a procurement process, without executive approval, and without committing to anything. At the end of 14 days they have data, specific to their environment, that they can take to their CISO and use to start a real conversation about stack optimization.

The message is not that the current stack is broken. The message is that every stack can perform better, and Pitwall shows you exactly how.

SECTION 3: TWO VOICE MODES

Pitwall has two distinct voices. They serve different purposes and they do not appear in the same document. Understanding which mode to use, and when, is one of the most important decisions in creating Pitwall content.

Declaration Mode

Declaration mode is for opening doors. It creates resonance. It names the situation the buyer is already living with, out loud, in language they recognize. It does not pitch. It does not feature-dump. It makes the reader feel understood.

Declaration mode is written peer to peer. It is the voice of someone who has been in the room, felt the frustration, and is speaking honestly about what they saw. It can use analogy and metaphor. Short sentences work well here. It should never feel like a vendor talking.

> *Declaration mode is used to create resonance, not to close deals. If a piece of content in declaration mode ends with a product claim, rewrite it.*

Declaration Mode: Reference Example

DECLARATION MODE

You built your security stack the right way. You researched the options, ran the POCs, and put in the work to deploy each product correctly. Over and over again. Each control added because a real threat demanded it.

But over time, as the threats grew and the stack grew with them, something changed. There were too many controls generating too much signal for anyone to keep up with. Your team started spending more time chasing alerts that turned out to be nothing than working the investigations that actually mattered.

You looked at AI SOC agents to help with triage. Maybe you deployed one. It helped. But in the back of your mind you know the underlying problem is still there. The stack should be working better. If you just had the time and the visibility to tune it properly, the whole operation would run differently.

You did not make wrong decisions. Your team is not missing anything. You just need to see how the car is running. That is what Pitwall gives you.

Evaluation Mode

Evaluation mode is for closing doors. Once a buyer is leaning in, the narrative gets out of the way. Evaluation mode is operator language: specific, outcome-oriented, and built for the person who needs to justify a decision internally.

No analogies. No emotional framing. No grand statements. Just the facts a skeptical SOC manager or CISO needs to trust that the product does what it says. Lead with outcomes, follow with specifics, let the numbers do the work.

> *Evaluation mode is used to support a decision already being considered. If a piece of content in evaluation mode opens with a problem statement or an analogy, rewrite it.*

Evaluation Mode: Reference Example

EVALUATION MODE

Pitwall gives SOC teams real-time visibility into how each security control is performing so they can identify and correct any control generating excessive false positives or duplicate alerts.

Pitwall automatically logs every alert generated across the security stack, from generation to resolution. Within approximately one week, it begins identifying which controls are operating outside accepted performance standards. From the dashboard, any underperforming control can be selected for auto-diagnosis. Pitwall generates human-readable output explaining the likely cause and specific recommended configuration changes. The SOC manager decides whether to implement the changes manually or enable Pitwall to apply them automatically.

Over time, Pitwall learns the environment and the team's response patterns. Recommendations tighten to reflect how the specific stack behaves and how the team likes to work.

Early adopters are seeing false positive reductions of up to 45% within 30 days. That reduction directly increases the time analysts have available for real investigations and backlogged tasks that have been deferred due to alert volume.

Voice Mode Assignment Rules

Content Type	Voice Mode	Notes
Homepage	**DECLARATION**	Opens with the problem named clearly. Three lane entry points lead to appropriate next pages. No product features on this page.
The Problem page	**DECLARATION**	The deepest declaration content on the site. No product mentions. Just the situation described honestly.
Product pages	**EVALUATION**	Feature and outcome focused. Operator language throughout. Numbers where available.
Platform datasheet	**EVALUATION**	One page. Buyer pain first, product specifics second, proof point last.
Shakedown landing page	**EVALUATION**	Specific offer, specific terms, specific outcome. No narrative warm-up needed here.

First call deck	**BOTH**	Opens with one declaration slide to create resonance. Shifts entirely to evaluation for the rest of the conversation.
Email outreach	**DECLARATION**	Short. Names the situation. Soft close. No product pitch in the first touch.
Case studies	**EVALUATION**	Customer voice, specific metrics, specific environment context. No superlatives.
LinkedIn posts	**DECLARATION**	Peer to peer. Observation or insight first. Never a product pitch.
Press releases	**EVALUATION**	Factual, specific, no hype language. Quotes from founders in operator voice.
Partner brief (external)	**EVALUATION**	Service revenue story. Specific commercial terms. No internal strategy detail.

SECTION 4: THE RULES

These rules apply to every piece of Pitwall content, internal or external, regardless of format or channel. They are not guidelines. If a piece of content violates any of these rules it does not go out.

Never Do These Things

- Never blame the buyer. Never imply the CISO made bad purchasing decisions. Never suggest the SOC team is not working hard enough. Never position the problem as something they should have caught or fixed sooner. The stack grew because the threat landscape grew. That is the only story we tell about how we got here.
- Never exaggerate or overclaim. No product delivers zero false positives. No stack becomes perfect. We have a 45% reduction from an early adopter. That is real and we use it with attribution. Everything else gets measured, specific language.
- Never use cybersecurity vendor jargon. Industry-leading, best-in-class, holistic security posture, single pane of glass, robust, comprehensive, cutting-edge. If a SOC manager would roll their eyes, it does not go in.
- Never talk down to the audience. They have been doing this longer than most vendors have existed. Write for someone who knows more about running a SOC than you do.
- Never use declaration and evaluation mode in the same document. The only exception is the first call deck, which opens with one declaration slide before shifting entirely to evaluation.
- Never make the F1 analogy a gimmick. It is a communication tool used to explain the problem and the product. It is not a brand personality. It does not appear in every document. When it fits, use it. When it feels forced, cut it.

Always Do These Things

- Lead with the outcome, not the feature. Your analysts get time back. Your stack performs the way it was designed to. Say that before you say how Pitwall makes it happen.
- Be specific. A 45% reduction in false positives in 30 days is more persuasive than any adjective. When you have a number, use it. When you do not, describe the mechanism in concrete terms.
- Write like an operator. Short sentences. Active voice. No hedging. If a sentence has more than three clauses, break it up.

- Respect the reader's intelligence. Do not over-explain. Do not state the obvious. Assume the person reading this has been in security for a decade and has heard every vendor pitch.
- Let the product earn the trust. The Shakedown exists because we believe the product speaks for itself when given the chance. The content's job is to get someone to the Shakedown. The Shakedown's job is to close the deal.

SECTION 5: APPROVED AND PROHIBITED LANGUAGE

The lists below are not exhaustive. They are calibration tools. When in doubt, test against the brand personality statement: does this sound like a trusted colleague, or does it sound like a vendor brochure?

Approved Language

USE THIS	INSTEAD OF THIS
Security stack	Security posture / security ecosystem
False positives	Noise (acceptable) / Alert fatigue (acceptable in context)
Tune / optimize	Remediate / resolve / fix (too clinical)
Real-time visibility	Comprehensive visibility / full visibility
Auto-diagnose	AI-powered remediation / intelligent automation
SOC manager / SOC team	Security operations professionals / practitioners
Get more out of your stack	Maximize ROI / optimize your investment
What is actually happening	Holistic view / 360-degree visibility
Your environment	Your infrastructure / your ecosystem
The Pitwall Shakedown	Free trial / proof of concept / pilot
Up to 45% reduction in false positives	Significant reduction / dramatic improvement
No infrastructure changes required	Agentless / frictionless deployment

Prohibited Words and Phrases

The following words and phrases are banned from all Pitwall content. No exceptions.

Industry-leading	**Best-in-class**	**Cutting-edge**
Holistic	**Robust**	**Comprehensive**
Single pane of glass	**Next-generation**	**Game-changing**
Empower	**Leverage**	**Seamless**
Transformative	**Unprecedented**	**Revolutionary**