The AI & Cyber Reference Guide

Part 1 — The No-Jargon Glossary

The words that come up in the room, defined for decisions.

AI awareness
Knowing AI exists and is relevant. Enough to follow the conversation but not enough to be accountable for the outcome.

AI literacy
Being able to ask the right questions, judge the claims, and own the decision. This is what leadership requires.

Narrow AI
A single-purpose tool doing one defined job (i.e., a spam filter, a fraud flag, spell-check). Everywhere already, reliable decisions.

Generative AI
Systems that create text, images, or code on demand. The tools your teams have often quietly adopted on their own.

Agentic AI
AI that's given a goal and acts on it, where the AI is planning and carrying out steps with limited human sign-off. The shift from assistant to actor.

Phishing
A fraudulent message engineered to trick a person into clicking, paying, or handing over credentials. Targets people, not firewalls.

Ransomware
Malicious software that locks up your data and demands payment to release it. Recovery speed matters more than prevention alone.

Business email compromise (BEC)
An attacker impersonates a trusted person (often an executive or vendor) to authorize a fraudulent payment or data transfer.

Supply-chain attack
A breach that reaches you through a vendor, partner, or piece of software you trust — now among the most common attack routes.

Human attack surface
Every point where a person's behavior can open the door: a click, a reused password, an over-shared access grant.

Governance (vs. management)
Management builds and runs the system. Governance sets policy, defines accountability, and owns the risk. You can delegate the first, never the second.

Risk appetite
How much risk your organization is willing to accept in pursuit of a goal. A leadership decision.

Incident response plan
A documented, rehearsed plan for what happens when something goes wrong: who acts, in what order, saying what to whom.

Exit plan
The answer to "can we turn this off, and what happens if we do?" No exit plan means the system has quietly become irreplaceable.

Part 2 — The Three Decisions Framework

One framework for every AI or security question.
The questions stay the same; your answers change with your context.

DECISION 1

Adopt or Assess?

Should we bring this AI in — and on what terms?

- What data does it use? (Sources, sensitivity, who else can see it.)
- Who is accountable for its outputs? (Name a human.)
- What's the failure mode? (When it's wrong, how wrong is it and how badly?)
- What's the exit plan? (Can we turn it off, and what happens if we do?)

Remember: ***"My team already uses it" is not a governance posture.***

DECISION 2

Govern or Delegate?

What does leadership own versus hand to technical teams?

- Risk appetite: How much AI risk is acceptable here? (Leadership's call, not IT's.)
- Accountability structure: Who answers when the system causes harm?
- Oversight cadence: How often does leadership review what's running?

Remember: ***Command authority doesn't transfer when you authorize someone to act; it expands.***

DECISION 3

Respond or Prepare?

Are we ready for an incident — or just hoping it won't happen?

- Is there an incident response plan leadership has read?
- What does "prepared" look like versus "hoping"? (Tested plan, known roles, backups, a communication template.)
- Can you run the 10-minute preparedness test today?

Remember: ***Preparation is the difference between a bad day and a business-ending disaster.***

The quality of your governance is based on the quality of your questions.