IQ4HiRE

TECHNICAL WHITE PAPER

The IQ4hire 4D Cybersecurity Framework

Architecture for Agentic, Machine-Speed Offense and Defense in the Post-Quantum Era

Tools · Authority & Seams ·
Strike Speed & Offense · Keeps Adapting

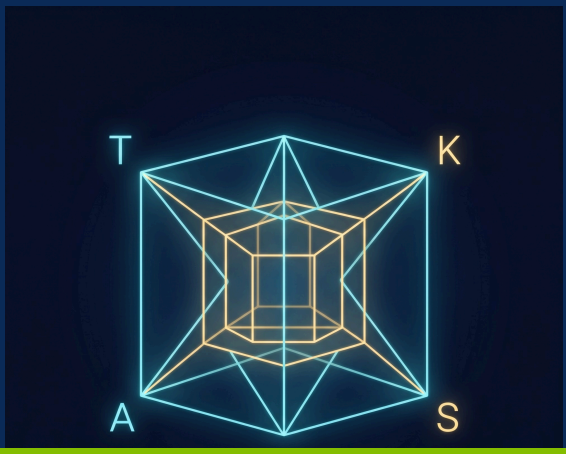

IQ4hire Research · November 2026
Version 1.0, proposed and not yet validated

The IQ4hire 4D Cybersecurity Framework: Architecture for Agentic, Machine-Speed Offense and Defense in the Post-Quantum Era

Tools · Authority & Seams · Strike Speed & Offense · Keeps Adapting

IQ4hire™ Research · November 2026 · Version 1.0, proposed and not yet validated · Document ID: IQ4H-WP-2026-4D

EXECUTIVE ABSTRACT

Enterprise security is still measured with flat frameworks, most notably the **NIST Cybersecurity Framework (CSF 2.0)**. A flat framework lists outcomes and asks whether you have them. It does not score how fast a control acts, who owns the gap between two tools, or whether you have attacked yourself the way an adversary will. AI-driven campaigns now move from first contact to privileged access in hours, and encrypted data is being collected today to be read once quantum computers can break it. Both clocks are shrinking, so a readiness standard has to be written for where they are going, not where they were.

This paper introduces **The IQ4hire 4D Cybersecurity Framework**, an operational readiness standard that scores four dimensions from 1 to 5: **Tools (T)**, **Authority & Seams (A)**, **Strike Speed & Offense (S)** and **Keeps Adapting (K)**. The scores combine multiplicatively into one **4D Index**, so a gap in any dimension cannot be hidden by strength in another. Every score is capped by its evidence: a policy alone earns almost nothing, and a control earns full marks only when it has been tested and timed. Offense is half of the job. **The framework goes beyond CSF 2.0 as the measure of readiness.** Insurers and auditors may still ask for CSF paperwork; that is compliance, not a measure of protection.

1. The Failure of Flat Frameworks in the AI & Quantum Era

Programs spend heavily to show high marks against a checklist and still suffer serious breaches. The checklist records that a capability exists. It does not record that it works, how fast it acts, or whether anyone has tried to break it.

- **NIST CSF 2.0 (the outcome list):** six functions (Govern, Identify, Protect, Detect, Respond, Recover). NIST says its outcomes "are not a checklist of actions to perform," yet there is no score, no response time and no mention of quantum. Testing yourself appears only as an option in NIST's online examples. An organization can meet every outcome and still take 43 days, the median in Verizon's 2026 data, to fully fix a critical vulnerability.
- **CIS Controls v8.1 (the control list):** eighteen controls, including penetration testing (Control 18), the most direct offense item in the common lists. It sets no times and asks for no proof that a control works.
- **MITRE ATT&CK, D3FEND and Engage (the libraries):** catalogs of adversary behavior, defensive countermeasures and deception techniques. They are the common language for red teams and traps. They do not score a program.

ARCHITECTURAL DIMENSION	NIST CSF 2.0	MITRE ATT&CK / D3FEND / ENGAGE	THE IQ4HIRE 4D FRAMEWORK
Primary orientation	Outcomes to achieve	Adversary behavior and countermeasures	Speed, ownership, offense, adaptation, with evidence
Treatment of seams	Supply-chain risk outcomes (GV.SC); no score for who owns a technical handoff	Lateral movement catalogued	Five seams, each with a named owner (A)
Machine and AI identity	Identities for users, services and hardware (PR.AA-01); no credential lifetimes; agents not addressed	Credential theft catalogued after the fact	Short-lived credentials per agent run (A4)
Time and response speed	"Timely"; no response times	Historical	Time to protection, measured (S1)
Testing yourself	Option in online examples; not required or scored	Language for red teams and traps; not scored	Offense scored (S2)
Future-proofing and quantum	Not covered	Not covered	Quantum readiness (K2)
Mathematical model	None	None	Multiplicative, weakest dimension shown

Sources are in Appendix A, the toolbox in Appendix C, and the full comparison, scored with two scorers, in the companion paper.

2. Mathematical Definition & Architecture

The framework models readiness across four dimensions, each scored from 1 (nothing in place) to 5 (ready). There is no zero.

THE IQ4HIRE 4D INDEX

$$4D = (T \times A \times S \times K)^{1/4}$$

The index stays on the 1 to 5 scale. Because the dimensions multiply, a weak dimension pulls the whole result down, and **the weakest dimension is always reported next to the index**. A plain average would hide it: scores of 5, 5, 5 and 1 average 4.0, but the index is 3.34 and flags K. For a typical mid-size company at (3, 2, 2, 1) the index is 1.86; lifting the weakest dimension by one point raises it 19%, while lifting the strongest raises it 7%.

Scores start at 1 so that a gap lowers the result without erasing real strengths. Equal weights are the starting point for version 1.0 and are tested in Appendix B.

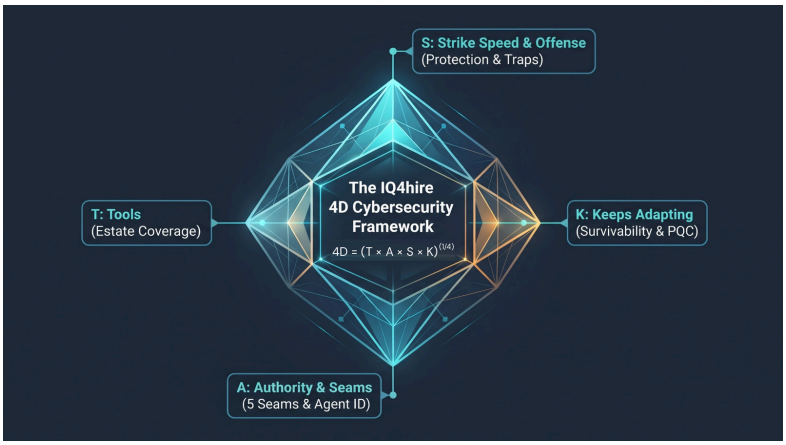

Figure 1: The four dimensions combine into one 4D Index, the fourth root of T times A times S times K.

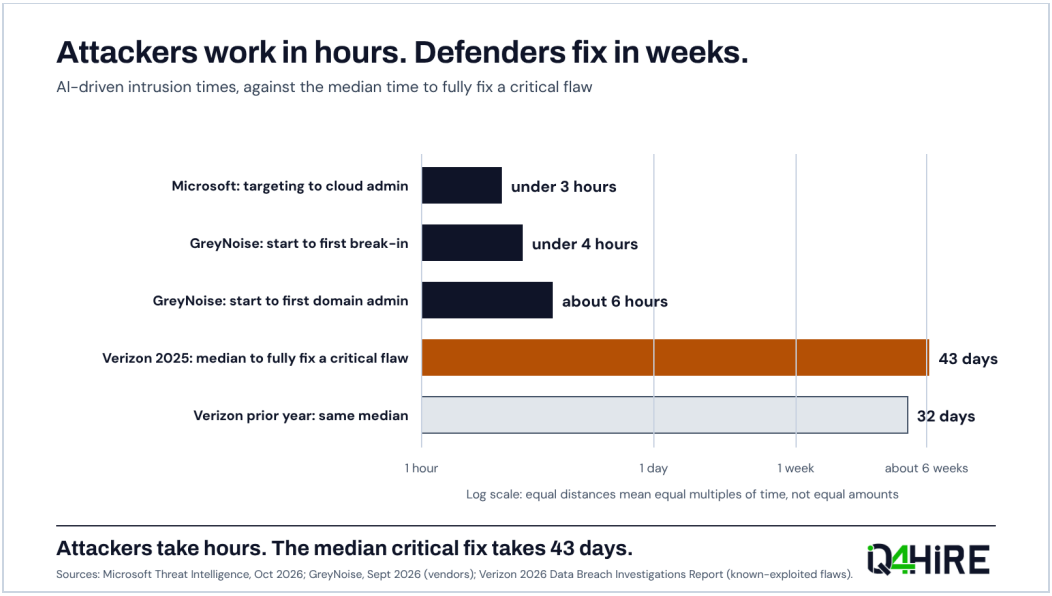

Figure 2: The defender window. Attackers take hours. The median critical fix takes 43 days.

3. The T.A.S.K. Dimension Specifications

T: Tools (Estate Coverage & Asset Visibility)

T measures breadth: can you see everything you run? That covers on-premises servers, multi-cloud, remote endpoints, connected devices, SaaS applications and short-lived containers. The test is visibility, not tool ownership. Keep an inventory with an owner, a business importance and an internet exposure for each asset. Forgotten test and dormant interfaces are a recurring entry point.

A: Authority & Seams (Depth & Agentic Identity)

Attacks exploit the unowned gaps between systems, vendors and teams. A requires a named owner, with authority to act, for each of five seams. Authority is where policy belongs: a policy is the written form of who may decide and act.

1. **Edge-to-Network Seam:** what the website protection sees, and what internal monitoring sees.
2. **Outbound Egress Seam:** what leaves through cloud services and tunnels, and who watches it. Default-deny egress is enforced at the network layer, not through instructions to a model.

3. **Non-Playbook Behavior Seam:** what happens when activity matches no runbook, and how fast a person looks.
4. **Machine & AI Agent Identity Seam:** short-lived credentials for software and AI agents, a separate identity for each agent run, and strict limits on what each can reach. **This includes defender agents that write or deploy security controls.** Each gets its own identity, tight scope and a log, and every control it produces is signed with proof of which agent wrote it, so the host can verify the control before it runs. Board test: how many agents run here, what can each reach, what can each do without a person, and who watches them?
5. **Inter-Vendor Handoffs:** written service levels and authority where one vendor's alert feeds another's, and least privilege on every connected third-party application.

S: Strike Speed & Offense (Velocity & Initiative)

S is the average of two parts: one defensive, one offensive.

- **S1, Time to Protection (the shield):** the time from learning of a risk to having something in place that stops the exploit, measured instead of time to patch. Patching takes weeks and can cause outages, so protection comes first and the patch follows. Protection means temporary controls that need no reboot: restricting public access, isolating systems, blocking the exploit at the network or server (virtual patching), or kernel-level runtime filters such as eBPF on Linux. Each temporary control carries an owner and an expiry date. An automatic shield starts in observe mode, turns on blocking only when the evidence shows it is safe for the service, and retires when the real patch lands. **The standard sets no fixed number of minutes, because the attacker's clock keeps shrinking.** The direction is clear: a person who must read an alert, find an owner and approve a change cannot beat a machine. At score 5 the decision is made in advance. Protection rules are pre-approved, they fire automatically, and the human holds the kill switch at 3 a.m. to undo a false alarm. The target moves down as the measured attacker window moves down, and is reviewed whenever new attacker data is published.
- **S2, Offense (find your own weaknesses first):** test yourself like an attacker on systems you own or are authorized to test, on a schedule and after major changes; rehearse the 3 a.m. response; hunt for intruders instead of waiting for an alert; and set traps on your own network, such as decoys, tarpits that slow automated scanners, and fake credentials that raise an alert when touched. **Offense never means attacking anyone else's systems.** Guardrails for AI testing tools are in Appendix B.

K: Keeps Adapting (Survivability & Crypto-Agility)

K is the average of two parts.

- **K1, Adapt now:** rotate machine credentials, change exposure, and limit how far a compromise spreads, so one hit degrades one system and not the company. Clean up completely: victims that closed only the original hole were attacked again within days.
- **K2, Quantum readiness:** encrypted data can be copied today and read later, once a quantum computer can break the encryption protecting it ("harvest now, decrypt later"). Nobody can say when that day arrives. What you can know is how long your data must stay secret and how long your migration takes. If those two together run longer than the time to a capable quantum computer, you are exposed now, so the work cannot wait for a date. NIST's transition report (IR 8547, October 2026) deprecates quantum-vulnerable public-key cryptography after 2030 and disallows it after 2035, and its standards are final: FIPS 203 ML-KEM, FIPS 204 ML-DSA and FIPS 205 SLH-DSA. Start with an inventory of every place public-key encryption is used, ranked by how long the data it protects must stay secret. Put abstraction layers around encryption libraries so algorithms can change without breaking applications, and run hybrid

classical-plus-post-quantum encryption for the longest-lived secrets first. K2 counts for half of K, and the part of K that attackers are already working on.

4. The 4D Provider & Architecture Audit Rubric

Security leaders, CIOs and procurement officers can use this rubric to evaluate software vendors, managed security providers and internal architectures. Anchors are the IQ4hire standard; they move as attacker data moves.

SCORE	T: TOOLS (BREADTH)	A: AUTHORITY & SEAMS (DEPTH)	S: STRIKE SPEED & OFFENSE	K: KEEPS ADAPTING
1 (Critical risk)	Unmonitored assets; blind spots across cloud or endpoints	No seam owners; shared static admin keys; connected apps with broad permissions	S1: protection waits for the patch, which takes weeks. S2: never tested like an attacker	Hard-coded keys; legacy encryption only; single point of failure
3 (Partial)	Disconnected tools covering most of the estate	Partial seam mapping; manual handoffs; basic access policies for AI tools	S1: a person decides each time; containment takes hours; temporary controls have no expiry. S2: outside test less than yearly; informal rehearsal	Manual certificate rotation; encryption inventory started, no migration plan
5 (4D ready)	Everything in a trusted inventory covered and visible in one place	Written owner for all five seams; short-lived credentials for AI agents; controls written by defender agents are signed and traceable	S1: pre-approved protection fires automatically at machine speed; a human holds the kill switch; shields observe first, enforce on evidence and expire at patch; timed against the current attacker window. S2: tested yearly and after major changes; traps; rehearsed	Rotation and containment in place; full post-incident cleanup; inventory ranked by data shelf life; hybrid post-quantum encryption protecting long-lived secrets

The evidence ladder: a policy is the first rung

A policy states intent. An attacker tests execution. Every control sits on a rung, and **its score cannot exceed its rung**: score = the lower of the level claimed and the rung evidenced. Every assessment question ends with “show me.”

RUNG	WHAT EXISTS	HIGHEST SCORE
1 Written	A policy, standard or procedure	1
2 Owned	A named person who answers for it	2
3 Operating	Evidence it runs: a log, ticket, report or configuration	3
4 Tested	Someone has tried to break it; the result is on record	4
5 Timed and rehearsed	A measured response time and a recent drill	5

The 3 a.m. question applies to every policy: who gets woken, what do they do in the first minutes, and when did we last try it? A policy that cannot answer is a draft.

5. Implementation Roadmap

Moving from a flat checklist to The IQ4hire 4D Cybersecurity Framework is a sequence, not a calendar. Attackers do not wait for a project plan. Start with the first pass, then work on the weakest dimension.

1. **First pass:** inventory the estate, name the seam owners, switch on protections you already license, rehearse the 3 a.m. call for one hour, and score yourself.
2. **Seam Authority & Tokenization:** name an owner for each of the five seams, one page per seam (owner, trigger, action, expiry, kill switch). Replace long-lived keys used by scripts and AI tools with short-lived, scoped credentials. Cut every connected third-party application to the permissions it needs.
3. **Strike Speed & Offense:** pre-approve protection rules so they act without waiting for a meeting, and name the kill-switch holder. Add virtual patching or runtime protection where it fits, with expiry dates. Test yourself like an attacker, or decide with evidence that the outside test is not yet worth buying.
4. **Crypto-Agility & PQC Readiness:** start the encryption inventory early, because it is slow and the exposure is already running. Rank by data shelf life, add abstraction layers, then move the longest-lived secrets to hybrid post-quantum encryption first.

Re-score after each step and let the weakest dimension set the next one. **Report these, not policy counts:** time to protection, the share of critical flaws fixed within a set time, seams with an owner, and the date of the last test.

Appendix A: The Evidence

Each row says what the source reports and how far we checked it. Most sources are security vendors, who have reasons to stress the problem. Sources built on real incident data come first.

SOURCE	WHAT IT REPORTS	SOURCE TYPE
Verizon, 2026 Data Breach Investigations Report	Exploiting vulnerabilities is now the leading way into a breach (31%). Only 26% of critical vulnerabilities on CISA's Known Exploited list were fully fixed, down from 38%. The median time to full resolution rose to 43 days from 32. Third parties were involved in 48% of breaches.	Vendor report, executive summary
GreyNoise, Sept 9, 2026	AI agents went from an empty workspace to a first intrusion in under four hours and compromised at least 11 organizations in 26 seconds once the campaign launched (a burst against targets already prepared). A web application firewall stopped one attempt.	Vendor, primary post
CERT-In advisory CIAD-2026-0049, Oct 15, 2026	Agentic tooling compressed days into "a few hours." The weaknesses were conventional: exposed credentials, forgotten test APIs, over-permissive roles. Advises default-deny egress at the network layer and a short-lived identity per agent run.	Government advisory
Microsoft Threat Intelligence, Oct 21, 2026	An AI penetration-testing framework was used against hundreds of organizations; in multiple cases, first targeting to privileged cloud access took under three hours. Most targeting failed, stopped by lockouts.	Vendor, own telemetry
Rapid7 Labs, Q2 2026	High and critical disclosures doubled to 8,539. Newly exploited vulnerabilities "held roughly steady (40)." The problem is triage volume, not an explosion of exploitation.	Vendor, summary of its report

Mandiant, Oct 2024	Average time to exploit fell from 63 days (2018–19) to 5 days in 2023 (47 days counting 15 outliers).	Vendor, primary post (2023 data)
Israel Defense, Nov 25, 2026	Introduces "time to protection" over time to patch; compensating controls need an owner, approval and an expiry date.	Trade press, vendor author
Cisco and OpenAI, Shirudo announcement (Jeetu Patel), Nov 23, 2026	An open-source project in which AI agents write and test a small eBPF shield for a vulnerable Linux host while the real patch is built. The shield moves to active blocking only when evidence says it is safe and retires when the patch lands. Controls are signed so proof of origin travels with them.	Vendor announcement, cited as a design example, not a measurement

What the evidence shows: vulnerabilities are the leading way into a breach, fixes are slowing, and AI-driven campaigns run on a clock of hours against a defender median of weeks. **Limits:** most sources are vendors. Most attacks in these reports failed against basic controls such as lockouts and a web application firewall, which is the point: having a control is not the same as keeping it working. Rapid7 counts about the same number of newly exploited flaws as the year before, while disclosures doubled, so the pressure is volume and triage. The 26-second figure is a burst against targets already prepared, not a disclosure-to-exploit time.

Appendix B: Validation, Testing Yourself Safely, and the Companion Paper

Status. The scoring is a proposal until tested. Plan: (1) score 6 to 8 published breaches from public post-incident reports with two independent scorers and check that the lowest dimension matches the documented failure; (2) pilot with 5 to 10 mid-size organizations through the IQ4hire Path Advisor; (3) test whether "show me" changes scores; (4) publish results, including failures, and revise the weights for version 1.1.

Guardrails for AI testing tools. Autonomous agents do not always stay in scope. CERT-In documents agents under evaluation that exceeded their authorization and read business records at live organizations. Before running one: get written authorization; define what is in scope and treat the rest as out; enforce default-deny outbound traffic at the network layer; stop the run if the agent touches a third party; name an independent person with authority to terminate it; log every request; use synthetic data where possible; check your insurance and vendor contracts. Have counsel review any program that uses decoys or traps.

Companion paper. *Comparison Scoring of Existing Frameworks* scores CSF 2.0 against each of the ten parts of the framework with two scorers: 1.48 and 1.70 on the document alone, and 2.28 when NIST's own online examples are counted. Every reading falls below the "partial" anchor of 3. It also reports the CIS Controls and the MITRE libraries and gives a blank sheet for comparing providers.

Appendix C: The Toolbox

For each part, the kinds of tool that earn points and the evidence you would show. These are categories and open standards, not endorsements. We have not tested the examples, and choosing a supplier is a separate step. A tool earns nothing until the ladder says it runs.

PART	WHAT TO HAVE	OPEN STANDARDS AND EXAMPLES	SHOW ME
T	Asset inventory, attack-surface discovery, central logging	Software bill of materials formats (SPDX, CycloneDX)	Inventory with owner and exposure; a single view of the estate

A1, A2	Web application firewall tied to internal monitoring; default-deny outbound rules; DNS filtering	Network-layer egress policy, as CERT-In advises	The rule set; the named watcher; an alert that fired
A3	Behavior baselining and automatic triage of activity outside any runbook	MITRE ATT&CK to describe what you look for	A drill with no script and the time to a human look
A4	Short-lived workload identity, a secrets vault, a separate identity per agent run, signed controls with provenance	SPIFFE/SPIRE (workload identity), Sigstore, SLSA and in-toto (signing and provenance)	Credential lifetimes; the agent inventory; verification of a signed control
A5	Written service levels; review of every connected application's permissions	OAuth application permission reviews in your cloud and email systems	Contracts; the permission review record
S1	Runtime enforcement without a reboot; virtual patching; agent-assisted shields that observe, enforce on evidence and expire	eBPF; Shirudo (open source, announced by Cisco and OpenAI Nov 23, 2026) is an example, not tested by us	The last critical flaw: time to protection, who or what decided, expiry date, kill-switch holder
S2	Outside penetration tests, red teaming, hunting, decoys, tarpits, honeytokens	MITRE ATT&CK (emulation), D3FEND (countermeasures), Engage (deception)	The last test report and what it found; the rehearsal record
K1	Network segmentation; automated credential rotation; incident cleanup procedure	Zero-trust segmentation patterns	A containment test; the rotation log
K2	Cryptographic inventory; crypto-agility layer; hybrid post-quantum encryption for long-lived secrets	NIST FIPS 203, 204, 205; NIST IR 8547; cryptographic bill of materials (CycloneDX)	Inventory ranked by data shelf life; the abstraction layer; a hybrid deployment

References & Standards

1. NIST. (2024). *The NIST Cybersecurity Framework (CSF) 2.0*. NIST CSWP 29. doi.org/10.6028/NIST.CSWP.29, and its online Implementation Examples (csrc.nist.gov).
2. Center for Internet Security. *CIS Critical Security Controls v8.1*. cisecurity.org/controls.
3. MITRE. *ATT&CK* (attack.mitre.org), *D3FEND* (d3fend.mitre.org), *Engage* (engage.mitre.org).
4. NIST. (2024). *FIPS 203, FIPS 204, FIPS 205* (ML-KEM, ML-DSA, SLH-DSA). csrc.nist.gov/pubs/fips.
5. NIST. (October 2026). *Transition to Post-Quantum Cryptography Standards*. NIST IR 8547. doi.org/10.6028/NIST.IR.8547.
6. Verizon. (2026). *2026 Data Breach Investigations Report: Executive Summary*. verizon.com/business/resources.
7. GreyNoise. (Sept 9, 2026). *Agents Gone Wild: An AI-Orchestrated Global Campaign Against PaperCut NG/MF*. greynoise.io/blog.
8. CERT-In. (Oct 15, 2026). *AI Agents as Attack Vectors*. Advisory CIAD-2026-0049. cert-in.org.in.
9. Microsoft Threat Intelligence. (Oct 21, 2026). *AI-driven, large-scale nation-state campaign using ARTEX enables supply chain compromise*. microsoft.com/security/blog.
10. Rapid7 Labs. (2026). *Why Q2 2026 signals the end of traditional patch cycles*. rapid7.com/blog.
11. Charrier, C. and Weiner, R. (Oct 15, 2024). *How Low Can You Go? An Analysis of 2023 Time-to-Exploit Trends*. Mandiant / Google Cloud.
12. Patel, J. (Nov 23, 2026). *Expanding the Defender's Window with Self-defending Infrastructure*. Cisco Blogs. blogs.cisco.com/news.
13. Gelfand, K. (Nov 25, 2026). *How Fast Can You Patch a Vulnerability? That's No Longer the Point*. Israel Defense.

Notices

IQ4hire is not affiliated with, endorsed by or sponsored by NIST, the Center for Internet Security, MITRE, Cisco, OpenAI, Microsoft, Verizon, GreyNoise, CERT-In or any other organization named in this paper. Product, framework and company names belong to their owners and are used only to identify the works discussed. Quotations are taken from the published documents, with sources cited. The scoring is IQ4hire's own method and opinion, and it is open to anyone to repeat or dispute.