CHAPTER 08 · THE ROAD AHEAD · ROADMAPS

Roadmaps and honest *timelines*.

Three eras — NISQ, early fault-tolerant, full FTQC. Four vendors with dates on slides. One resource estimator that tells you which dates are physics and which are marketing. By the end of this chapter you'll read any quantum roadmap with the right kind of scepticism.

WHAT YOU'LL LEARN

- Place every vendor on the NISQ → EFT → FTQC line
- Run Q# resource estimation on Shor
- Tell physics from marketing

THE MAP · WHERE WE ACTUALLY ARE

NISQ → *early fault-tolerant* → FTQC.

2016 – 2023 · NISQ

Noisy, no QEC

Hundreds of physical qubits. No error correction. Circuit depth bounded by raw two-qubit fidelity. Useful science narrow — RCS benchmarks, small VQE, advantage debates.

2024 – 2028 · EFT

Below-threshold

Tens of logical qubits. QEC demonstrated on hardware — Willow (Dec 2024), Quantinuum Helios (Nov 2025, 48 logical), IBM Kookaburra targeting qLDPC memory. Useful science widening but narrow.

2029+ · FTQC

Useful, deep, scaled

Thousands of logical qubits. Deep circuits. Targets: useful Shor, large chemistry, materials. IBM Starling, Google 2029 target, Quantinuum 2030 universal — all conditional.

As of May 2026 the honest answer is: *early EFT*. Below-threshold demonstrations exist; useful FTQC does not. Every claim past that line is a roadmap, not a result.

— VENDOR ROADMAP · IBM

Heron → Loon → Kookaburra → *Starling*.

2024 · HERON R2

156 qubits

Production today. Mid-circuit measurement, dynamic circuits, tunable couplers. The current commercial baseline.

2025 · LOON

c-couplers

Cross-chip long-range connectivity between modules — the prerequisite for leaving the surface code behind.

2026 · KOOKABURRA

qLDPC memory

First module targeting Gross code [[144,12,12]] — 12 logical qubits from 144 physical. An order of magnitude in overhead vs surface code.

2028 – 2029 · STARLING

~200 logical

IBM's stated first useful FTQC machine. ~10^4 physical qubits supporting 10^8 gate operations on logical qubits.

Heron is shipped silicon. Loon and Kookaburra are 2025-2026 deliveries. *Starling is a 2028-2029 target conditional on engineering nobody has yet delivered.* Read every roadmap node with that calibration.

— VENDOR ROADMAPS · THE OTHER THREE

Google, Quantinuum, Microsoft — three timelines, *three confidence levels*.

GOOGLE · 2029 TARGET

Useful FTQC

Public target — a useful error-corrected processor by 2029. Same horizon as IBM Starling. Backed by Willow (Dec 2024) below-threshold result on a distance-7 surface code. *Conservative.*

QUANTINUUM · 2030 TARGET

Universal FTQC

Roadmap puts universal fault-tolerant QC at 2030. Helios shipped Nov 2025 — 98 barium ions, 99.92% two-qubit fidelity, 48 logical qubits as a bridge. *Conservative.*

MICROSOFT · "YEARS NOT DECADES"

1M qubits

Majorana 1 (Feb 2025) — claimed path to a million qubits on a chip via topological qubits. Per chapter 6 the underlying physics is contested. *Treat sceptically — needs corroboration before you plan around it.*

Two timelines (Google, Quantinuum) line up with the physics and the published demonstrations. One (Microsoft) is the kind of claim where the only honest answer in May 2026 is *wait for independent replication*.

WORKED EXAMPLE · SHOR AT RSA-2048

How big does the machine have to be? *Three numbers.*

~4,000

LOGICAL QUBITS

Shor at RSA-2048 needs roughly four thousand error-corrected qubits. Today's best: ~50 logical.

$\sim 10^8$

T-GATES

A hundred million non-Clifford T-gates per factoring run. Each one needs a magic-state factory upstream.

~20M

PHYSICAL QUBITS

Surface code at distance 25, including factories. Three orders of magnitude beyond any 2026 machine.

Even under optimistic assumptions, useful Shor is *5 to 10 years* away in May 2026. Roadmaps that promise sooner are either talking about smaller problems (RSA-1024, ECC-256) or skipping a step in the overhead math. The next slide shows you the tool that does this math out loud.

FRAMEWORK ANCHOR · MICROSOFT Q# + AZURE QUANTUM

The resource estimator, in *eleven lines*.

QSHARP · AZURE-QUANTUM · PYTHON

```
import qsharp
from qsharp.estimator import EstimatorParams, QubitParams, QECScheme

qsharp.eval('open Microsoft.Quantum.Crypto.Shor;')

params = EstimatorParams()
params.error_budget = 1e-3
params.qubit_params.name = QubitParams.GATE_NS_E3      # superconducting, 1 ns gate, 1e-3 error
params.qec_scheme.name  = QECScheme.SURFACE_CODE

result = qsharp.estimate('FactorInteger(2048)', params=params)
print(result.summary)   # logical qubits, runtime, physical qubits, T-factories
```

Real syntax — runs against the public Azure Quantum Resource Estimator. Swap `SURFACE_CODE` for a qLDPC scheme, or the qubit params for trapped-ion presets, and the answer changes. *When a vendor promises useful Shor by year X, run this against year X assumed hardware and check the math closes.*

— THREE NUMBERS VENDORS QUOTE · AND WHAT EACH ONE HIDES

Quantum Volume, CLOPS, AQ — *none tells the full story.*

QV

QUANTUM VOLUME · IBM

Single number combining qubit count, connectivity, fidelity, depth. Doubles each Heron release. *Flatters connectivity* — heavy-hex IBM layouts score high; grid layouts lower.

CLOPS

CIRCUIT-LAYER OPS PER SEC

Throughput — how fast the stack pushes circuit layers through control electronics + Runtime. *Flatters throughput*, says nothing about fidelity per layer.

AQ

ALGORITHMIC QUBITS · IONQ

Counts how many qubits a real algorithm can use end-to-end at acceptable fidelity. *Flatters trapped-ion architectures* — all-to-all connectivity inflates AQ vs nearest-neighbor.

Each metric is honest in isolation; *none is sufficient alone*. Triangulate: T_1, T_2, two-qubit fidelity (the chapter-4 baseline) *plus* one aggregate number. Beware any roadmap chart that omits the baseline.

— EXERCISE · 30 MINUTES · FREE AZURE QUANTUM TIER

Estimate three problems. Plot one chart. Read a roadmap like a *physicist*.

— STEP 1 · INSTALL

```
pip install qsharp azure-quantum
```

Free tier on Azure Quantum gives unlimited estimator calls — the estimator runs locally, no QPU time billed.

— STEP 2 · THREE RUNS

Run **`qsharp.estimate`** three times:

- `FactorInteger(1024)` — RSA-1024
- `FactorInteger(2048)` — RSA-2048
- `EllipticCurveDLog(256)` — ECC-256

— STEP 3 · PLOT + INTERPRET

Log-scale plot of physical qubit count vs the three problems. Then swap **`SURFACE_CODE`** for the Gross qLDPC code in **`params.qec_scheme`** and re-run. Watch the answer drop by roughly an order of magnitude — *that is the IBM Kookaburra bet, quantified.*