PART I - WHY MICROSERVICES ARE NOT THE STARTING POINT

1. The Microservice Temptation

The first microservice mistake is not technical. It is psychological. A system becomes difficult, and a team begins to believe that deployment boundaries will fix modelling problems.

Every developer has felt the attraction. The codebase is large. The solution has many projects. The database has too many tables. A feature crosses AlarmManagement, RootCause, Audit, Compliance, and Reporting. Someone says the sentence that sounds practical and modern at the same time: let us split it into microservices.

The sentence is dangerous because it can mean two completely different things. It can mean: we have mature ownership boundaries and independent release pressure. Or it can mean: the application is hard to understand, so we want the network to make the boundaries visible. The second meaning creates pain, not architecture.

> **PLAIN EXPLANATION**
> A microservice is not a small Web API. It is a responsibility boundary that owns data, language, contracts, failure behaviour, and operations. If those responsibilities are unclear before extraction, distribution will make them harder to see, not easier.

Chapter 1 has one job: slow the decision down. It does not argue against microservices. It argues against beginning with them.

What you already know

You already know how a clean modular application can be organized. You may have a Domain project, an Application project, an Infrastructure project, and an API project. You may already separate commands from queries. You may already keep EF Core mapping out of your entities. You may already publish domain events inside the application boundary.

That structure is good. It is also not a microservice architecture. It is a modular architecture. The difference matters.

STRUCTURE	WHAT IT PROVES	WHAT IT DOES NOT PROVE
Project boundary	The code can be organized and dependencies can be controlled.	That the project deserves a separate runtime.
Bounded context	A language and model boundary exists.	That the context needs independent deployment.
SQL schema	Data ownership can be made visible in the database.	That the database should be split immediately.
Web API controller	A client can call a capability.	That the controller is a service boundary.
Integration event	A fact can be published outside the model.	That every reaction should be asynchronous.

The temptation begins when we confuse these proofs. A bounded context proves a language boundary. A service requires more: independent data, stable contracts, failure isolation, observability, deployment reason, and a team that can own the runtime consequences.

> **NEXUS-1 REMINDER**
> AlarmManagement, RootCause, Audit, Compliance, Reporting, DigitalTwin, Instrumentation, and ReactorFleet are meaningful sectors. Meaningful sectors are service candidates, not automatic services.

The naive split

The naive split takes the visible folders and turns them into deployables. AlarmManagement becomes an Alarm service. RootCause becomes a RootCause service. Compliance becomes a Compliance service. Audit becomes an Audit service. Reporting becomes a Reporting service. At first the diagram looks cleaner because each box has a name.

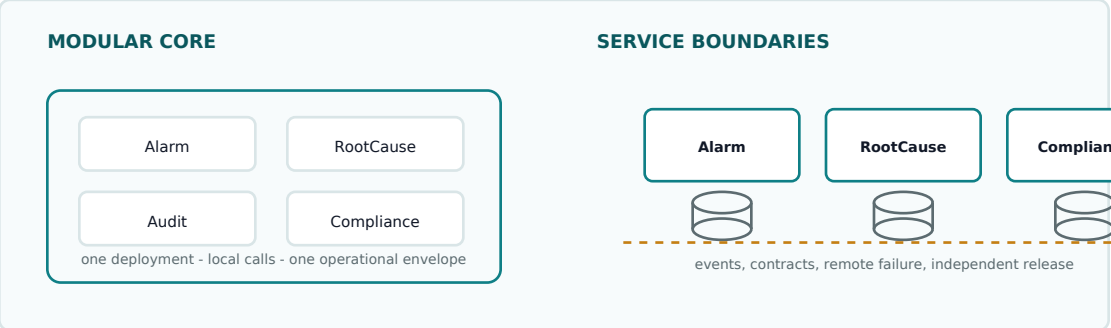

The diagram has become more impressive, but the system has not necessarily become more coherent. A local method call has become a remote call. A database transaction has become a consistency question. A class reference has become a public contract. A log line has become a distributed trace. A unit test has become a contract test. A deployment has become an operational event.

> **COMMON MISTAKE**
> Do not mistake the visual neatness of a distributed diagram for architectural maturity. A distributed monolith often looks excellent on a slide and behaves badly in production.

The microservice temptation is powerful because it promises clarity without first doing the slow work that clarity requires: language ownership, data ownership, contract ownership, and failure ownership.

Why the naive solution fails

When a modular flow is pushed across service boundaries too early, normal software events become distributed failure modes. Nothing supernatural happens. The system simply starts charging interest on every unclear boundary.

NAIVE MOVE	IMMEDIATE RESULT	HIDDEN COST
Put RootCause behind an HTTP API	Other services can call it directly.	Every caller now depends on RootCause availability and contract stability.
Let services share the same database	Migration looks easy.	No service truly owns its data; the system remains coupled at the most dangerous layer.
Publish every domain event publicly	Integration looks rich.	Private model language becomes public contract and is difficult to change.
Make Reporting query operational tables	Screens are easy to build.	Reporting becomes a hidden dependency on source-of-truth schemas.
Retry every failed call	Failures appear recoverable.	Retry storms can amplify an outage across the system.

The cost is not only latency

Latency is the easiest cost to notice. It is not the most important. The deeper cost is the loss of local certainty. Inside one transaction, the system either commits the root-cause verdict and its local state or it does not. Across services, the verdict may be committed, the event may be waiting in the outbox, Audit may be temporarily unavailable, Compliance may not yet know, and Reporting may still show the old read model.

> **ARCHITECTURAL PRINCIPLE**
> A distributed system must be designed around partial progress. If the design assumes that all services move together, it is not a microservice architecture. It is a distributed transaction wish.

NEXUS-1: a context map is not a deployment map

NEXUS-1 already has strong names: ReactorFleet, Instrumentation, AlarmManagement, RootCause, DigitalTwin, ReinforcementLearning, Audit, Compliance, Reporting, and the operator console. The names are valuable because they express different responsibilities and language areas. But the context map answers a different question from the deployment map.

MAP TYPE	QUESTION IT ANSWERS	NEXUS-1 EXAMPLE
Context map	Where does language change?	AlarmEvent is not Evidence; Recommendation is not Command.
Data map	Who owns the source of truth?	RootCause owns RootCauseCase and its verdict. Reporting owns projections, not truth.
Contract map	What can cross the boundary?	RootCauseVerdictIssued.v1 is public; VerdictIssued remains internal.
Deployment map	What can be released and operated independently?	RootCause may become a service only if the operational cost is justified.

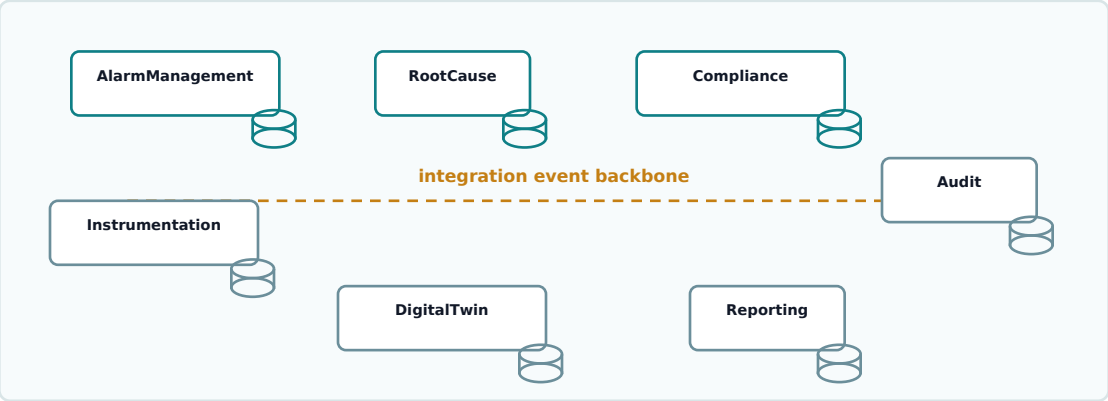

Figure 1.1 - NEXUS-1 sectors as service candidates. A candidate map is an evaluation surface, not a promise that every box will become a separate deployment.

TERMINOLOGY DISCIPLINE

Use RootCauseVerdictIssued.v1 for the public integration event. Keep VerdictIssued as the internal domain event. Do not replace this with RootCauseCaseClosed or RootCauseVerdictPublished, because those names change the language of the decision.

What a service must own

The smallest useful definition of a microservice is not about size. It is about ownership. A service is a deployable unit that owns a business capability, protects its internal model, controls its data, exposes stable contracts, tolerates remote failure, and can be observed when something goes wrong.

OWNERSHIP	MEANING	QUESTION TO ASK
Language	The service controls what its words mean inside its boundary.	Does RootCause define Evidence, Hypothesis, and Verdict without borrowing UI or Reporting terms?
Data	The service owns the source-of-truth tables for its capability.	Can another service change a RootCauseCase directly? The answer should be no.
Decision	The service decides only what it is responsible for deciding.	Does Compliance decide compliance review, while RootCause decides root cause?
Contract	The service publishes APIs and events that do not leak private aggregate shape.	Is RootCauseVerdictIssuedV1 stable even if RootCauseCase internals change?
Failure	The service has explicit behaviour when dependencies are unavailable.	What happens if Audit or Compliance is down when a verdict is issued?
Operations	The service can be monitored, deployed, configured, and diagnosed.	Can we trace a case from alarm flood to compliance review?

Only after these ownerships are credible does the microservice conversation become honest. Before that point, the safer design is usually a modular core with strong internal boundaries and a small number of well-governed integration seams.

SERVICE READINESS TEST

A service is ready when deleting the network boundary would not be the only thing keeping its model understandable. The boundary must already exist in language, data, and rules.

A small C# shape: the tempting synchronous chain

The following example is intentionally simple. It shows the style of code that often appears when a modular flow is split into services too quickly. The handler looks readable, but it hides a distributed transaction expectation inside ordinary sequential code.

```
public sealed class AlarmFloodHandler
{
    private readonly IRootCauseClient _rootCause;
    private readonly IAuditClient _audit;
    private readonly IComplianceClient _compliance;

    public async Task Handle(AlarmFloodDetectedV1 flood)
    {
        var caseId = await _rootCause.OpenCaseAsync(flood);
        var verdict = await _rootCause.AttachEvidenceAndDecideAsync(caseId);

        await _audit.WriteAsync(verdict);
        await _compliance.RequestReviewAsync(verdict);
    }
}
```

Nothing in this code says what happens if Audit is unavailable after the verdict has been committed. Nothing says whether Compliance may receive the review request twice. Nothing says whether the operator console should show a verdict before Reporting catches up. Nothing says which correlation ID ties the calls together. The code is clean; the architecture is not finished.

> **BETTER READING OF THE SAME FLOW**
> The verdict should be committed by the service that owns it. A public integration event should be recorded in an outbox in the same local transaction. Other services should react idempotently, with correlation, retries, and observable failure handling.

This does not mean synchronous calls are forbidden. It means synchronous calls must be chosen deliberately. Some queries are acceptable. Some commands are acceptable. Chained decisions that pretend to be one transaction are the trap.

The better first question

The first question is not how many services should we create? The better first question is: which boundaries are already true even if we do not deploy them separately?

BOUNDARY SIGNAL	WHAT IT LOOKS LIKE IN NEXUS-1
Different language	Evidence, Hypothesis, Verdict, Recommendation, AuditEntry, and ComplianceFinding are not synonyms.
Different source of truth	RootCause owns the verdict; Reporting owns a projection of it; Audit owns evidence of change.
Different reason to change	Alarm flood detection evolves differently from compliance review policy.
Different failure tolerance	Reporting may lag; RootCause verdict issuance must remain locally consistent.
Different operational load	Alarm events and signal readings may have different volume and retention needs from compliance records.
Different public contract	RootCauseVerdictIssued.v1 is a stable public fact; RootCauseCase internals are private.

When several signals point to the same boundary, the service candidate becomes credible. When only the folder structure points to the boundary, extraction is premature.

DESIGN ATTITUDE

A mature architecture is allowed to say no. Keeping a context inside a modular monolith can be the disciplined choice when independent deployment would add more risk than value.

The first NEXUS-1 service-candidate reading

Chapter 1 does not decide the final service architecture. It only establishes the first reading: some sectors are strong candidates, some are conditional, and some should remain inside a protected modular core until later evidence justifies extraction.

CATEGORY	CANDIDATES	EARLY REASONING
Strong candidates	AlarmManagement, RootCause, Compliance, Audit, Reporting	They have clearer language, data, public events or read models, and plausible independent evolution.
Conditional candidates	DigitalTwin, Instrumentation, ReinforcementLearning Advisory	They may benefit from separation, but coupling, volume, timing, and advisory/control boundaries must be evaluated carefully.
Possibly kept together	ReactorFleet, simulation internals, UI composition, shared demonstration infrastructure	These may be too coupled, too foundational, or too demonstration-specific to split early.
Never implied by this book	Safety-like control logic for a real plant	NEXUS-1 is educational and advisory. It does not claim operational plant authority.

The point is not to be timid. The point is to be exact. Strong boundaries should be made stronger. Weak boundaries should be studied before they are distributed. Artificial boundaries should be removed, not wrapped in HTTP.

MICROSERVICE TEMPTATION RESOLVED

The promise of microservices is not smaller code. The promise is independent evolution of responsibly owned capabilities. When that independence is not real, a service is only another place where the same coupling can fail.

Chapter checkpoint

You are ready to continue when you can answer these questions without using tool names as explanations.

- Why is a bounded context not automatically a microservice?
- What ownerships must a service have before extraction is credible?
- Why can a shared database make a microservice architecture false even when the APIs are separate?
- What failure is hidden inside a simple synchronous service chain?
- Why is RootCauseVerdictIssued.v1 a public integration event while VerdictIssued remains internal?
- Which NEXUS-1 sectors look like strong candidates, and which require caution?
- What does the phrase split only when the boundary is already true mean in practical design work?

HONEST BOUNDARY
This chapter has not designed a broker, a database-per-service topology, an API gateway, deployment units, containers, or observability. Those come later. Chapter 1 only establishes the judgement rule: distribution must follow responsibility, not fashion.

Next chapter

Chapter 2 will separate four ideas that are often collapsed into one: module, bounded context, service candidate, and deployable microservice. The distinction is the foundation for every service boundary decision that follows.

Do not split a system because microservices are modern. Split only when the boundary is already true.

PART I - WHY MICROSERVICES ARE NOT THE STARTING POINT

2. Why a Bounded Context Is Not Automatically a Service

A bounded context is a language boundary first. A microservice is a runtime boundary as well. Confusing the two turns a good DDD idea into a distributed-system problem before the system is ready for it.

The previous chapter slowed down the microservice decision. This chapter makes the main distinction precise. In Domain-Driven Design, a bounded context tells us where a model has a controlled meaning. It says that words such as Alarm, Evidence, Verdict, Recommendation, AuditEntry, and ComplianceFinding must not be blended into one universal model. That is already valuable. But it is not yet a deployment decision.

A microservice adds more than a boundary of meaning. It adds a process boundary, a network boundary, a data-ownership boundary, a release boundary, and an operational boundary. Those additions are not free. They only become worthwhile when the context can carry them without leaking its model or depending on another service to remain correct.

PLAIN EXPLANATION
A bounded context answers: what does this word mean here, and who owns the rule? A microservice also answers: who owns the data, what is the public contract, what happens when the network fails, and who operates this thing at runtime?

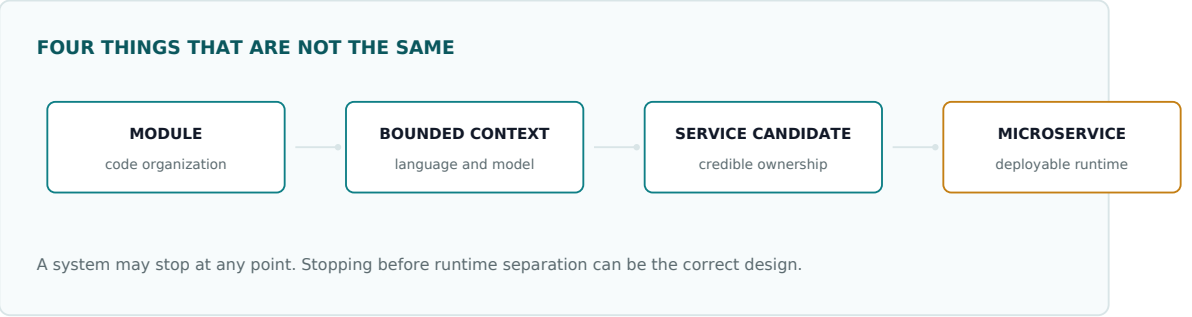

Figure 2.1 - Module, bounded context, service candidate, and microservice are related but not identical. A clean architecture can stop before runtime separation and still be correct.

What you already know

You already use boundaries before you call them microservices. A namespace is a boundary. A project is a boundary. A SQL schema is a boundary. A folder of EF Core configuration classes is a boundary. An aggregate is a boundary. A REST controller is a boundary. None of these, by itself, proves that the bounded thing should run as a separate service.

This is especially important in NEXUS-1 because the system already has many meaningful sectors. ReactorFleet, Instrumentation, AlarmManagement, RootCause, DigitalTwin, ReinforcementLearning, Audit, Compliance, and Reporting all have a reason to exist. If we treat existence as extraction, the service map becomes automatic. Architecture should not be automatic.

BOUNDARY TYPE	WHAT IT PROTECTS	WHAT IT DOES NOT SETTLE
Aggregate boundary	Local invariants and consistency rules.	Whether the aggregate deserves a service.
Bounded context	Language, model, and ownership of meaning.	Whether runtime separation is justified.
SQL schema	Database namespace and referential ownership.	Whether database-per-service is ready.
Project boundary	Compile-time dependencies and code organization.	Whether remote communication is needed.
API boundary	How a caller reaches a capability.	Whether the internal model is protected.

NEXUS-1 REMINDER

A sector is a candidate, not a verdict. Each candidate must be tested for language ownership, data ownership, contract stability, failure behaviour, and operational responsibility.

The four-step confusion

Many microservice failures begin with a quiet substitution. The team says bounded context, then draws a service. It says service, then creates a Web API. It says independent data, then keeps the same shared database. It says event-driven, then publishes whatever the aggregate currently looks like. Each step sounds close enough to the previous one that the mistake is easy to miss.

CONFUSED STATEMENT	WHY IT IS WRONG	BETTER STATEMENT
This is a bounded context, so it is a service.	A bounded context is a model boundary; a service is a runtime and operational boundary.	This context may become a service if the readiness tests pass.
This has its own controller, so it is a service.	A controller is an entry point, not an ownership model.	The controller must protect a capability owned behind it.
This has its own tables, so it owns data.	Tables in one shared database may still be changed by other modules.	Ownership means only the owning service writes source-of-truth data.
This publishes events, so it is loosely coupled.	Bad events leak private model shape and create brittle consumers.	Public integration events must be deliberate contracts.

The difference is not academic. In the alarm-to-compliance flow, a careless split can make every step appear independent while still requiring all services to be online, all databases to be consistent immediately, and all consumers to understand the private structure of RootCauseCase. That is not microservice architecture. It is a distributed monolith with better labels.

COMMON MISTAKE
Do not let the bounded-context diagram become a deployment diagram by habit. A context map is a map of language and relationships. A service topology is a map of runtime responsibility.

The service extraction ladder

A safer way to think is to place every context on a ladder. The bottom steps are DDD steps. The upper steps are distributed-system steps. A context may be valuable even if it never reaches the top.

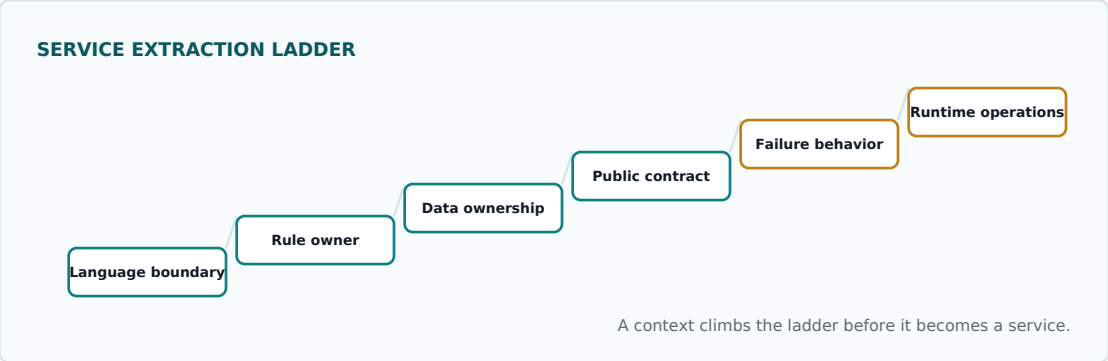

Figure 2.2 - The extraction ladder. Runtime separation comes after language, rule ownership, data ownership, public contract shape, failure behaviour, and operational ownership.

The ladder matters because it prevents a premature leap. A context that has a clear language but no independent data ownership should not become a service yet. A context that owns data but has no stable public contract should not expose its aggregate shape as an API. A context that publishes a stable contract but has no failure policy is still incomplete.

STEP	NEXUS-1 QUESTION
Language boundary	Does RootCause own Evidence, Hypothesis, Verdict, and Case without borrowing Reporting or UI words?
Rule owner	Which aggregate or domain service protects the invariant?
Data ownership	Can any other service write the source-of-truth rows directly?
Public contract	Is the contract stable enough for consumers that do not know the private aggregate?
Failure behaviour	What happens if Compliance, Audit, or Reporting is down?
Runtime operations	Can the service be deployed, monitored, traced, and diagnosed independently?

A NEXUS-1 example: RootCause

RootCause is a strong teaching example because it looks like an obvious service candidate. It has a language of its own. It owns a decision. It protects invariants. It publishes a meaningful public fact. It is also dangerous to extract casually, because its inputs and outputs sit inside a wider flow.

Inside the RootCause context, an internal domain event may be named VerdictIssued. That name is allowed to be local because it belongs to the model. Outside the context, the published integration event should be RootCauseVerdictIssued.v1, represented in C# by a public contract type such as RootCauseVerdictIssuedV1. The version number is not decoration. It is a promise that consumers can depend on.

```
// Internal model event: local language, not a public contract.
public sealed record VerdictIssued(
    RootCauseCaseId CaseId,
    VerdictText Verdict,
    DateTime OccurredAtUtc);

// Public integration event: stable contract for other services.
public sealed record RootCauseVerdictIssuedV1(
    Guid EventId,
    string CorrelationId,
    int RootCauseCaseId,
    int UnitId,
    string VerdictCode,
    string VerdictSummary,
    DateTime IssuedAtUtc);
```

The two events are related, but they are not the same object. The internal event can change when the aggregate changes. The public integration event must change only through a versioned contract. This is the kind of discipline a bounded context must show before it becomes a service.

> **TERMINOLOGY DISCIPLINE**
> AuditEntry is evidence, not workflow. A recommendation is not a command. A read model is not the source of truth. These sentences are not style preferences; they are service-boundary protection rules.

When a context should remain inside the modular core

The modular monolith is not a failure state. It is often the healthiest place for a context while its language, data, contracts, and failures are still being learned. A context can be well designed inside one deployable application. It can have its own namespace, project, schema, tests, domain events, and outbox table without being separately deployed yet.

KEEP IT MODULAR WHEN...	PRACTICAL MEANING
The boundary is mostly technical.	A folder exists, but no separate business capability is visible yet.
Data ownership is unclear.	Several modules still write the same source-of-truth rows.
Contracts are unstable.	Consumers need private aggregate details to do their work.
Failure behaviour is unknown.	The team cannot explain what happens when a dependency is unavailable.
Operational ownership is missing.	No one can own monitoring, deployment, incidents, and versioning for the service.
The flow requires immediate consistency.	Splitting would force distributed transactions or fragile synchronous chains.

For NEXUS-1, this means we should not rush every sector out of the core. Some sectors are shared foundations. Some are high-volume fact stores. Some are read-side projections. Some are advisory. Each has a different reason to exist and therefore a different extraction threshold.

HONEST DESIGN CHOICE
Keeping a context modular is not avoiding architecture. It can be the architecture that protects the model until the evidence for distribution is strong enough.

When a context can become a service candidate

A context becomes a credible service candidate when its ownership remains meaningful even after the method call becomes remote. The boundary must survive latency, retries, duplicates, versioning, partial availability, and independent deployment.

SERVICE-CANDIDATE SIGNAL	NEXUS-1 EXAMPLE
It owns a decision.	RootCause issues the verdict; Compliance decides review status; AlarmManagement detects the flood.
It owns source-of-truth data.	RootCause owns root-cause cases and verdicts, not Reporting.
It publishes stable facts.	RootCauseVerdictIssued.v1 and AlarmFloodDetected.v1 are public facts, not private entities serialized by accident.
It can tolerate asynchronous reaction.	Compliance review can be triggered after the verdict event without rolling back the verdict.
It has a different scaling or release pressure.	Reporting projections and alarm ingestion may evolve differently from compliance review.
It can be observed independently.	Correlation IDs connect AlarmManagement, RootCause, Audit, Compliance, and Reporting.

Notice the order. The signal is not that the code is large. The signal is that the responsibility is independent enough to justify the cost of remote communication and independent operation.

SERVICE CANDIDATE RULE
A context can become a service candidate when its independence is real before deployment. Deployment should reveal a boundary, not invent one.

Crossing the boundary without leaking the model

The strongest test of a service candidate is what it exposes. A weak service exposes tables, entity names, or aggregate internals. A stronger service exposes commands, queries, and events that describe public intent or public facts.

PRIVATE MODEL SHAPE	PUBLIC CONTRACT SHAPE
RootCauseCase.FinalVerdict	RootCauseVerdictIssued.v1
Hypothesis.Status changed to Promoted	RootCauseVerdictIssued.v1 after decision
AuditEntry inserted	No workflow event; audit records evidence
Recommendation row created	RecommendationAvailable.v1 or queryable read model, not an actuator command
ComplianceFinding table updated	ComplianceReviewRequested or ComplianceReviewCompleted, depending on the owned decision

This distinction protects future change. RootCause may later change its internal hypothesis scoring, evidence storage, or case lifecycle. Consumers should not break because they were coupled to private shape. They should depend on a public fact: a verdict was issued, by this service, for this case, at this time, under this correlation.

PUBLISHED LANGUAGE
A public contract is not a convenient DTO. It is the published language of the service boundary. It deserves versioning, examples, tests, and review.

A first readiness reading of NEXUS-1 sectors

The table below is not the final architecture. It is the first disciplined reading. It asks whether each sector looks closer to a module, a bounded context, a service candidate, or a deployable microservice. Later chapters will test these readings against data ownership, APIs, messaging, workflows, and operations.

SECTOR	CURRENT READING	REASON
AlarmManagement	Strong service candidate	Owns alarm definitions, alarm events, flood detection language, and the public AlarmFloodDetected.v1 fact.
RootCause	Strong service candidate	Owns cases, evidence, hypotheses, verdicts, and RootCauseVerdictIssued.v1.
Compliance	Strong but workflow-sensitive	Owns review policy and findings, but must not confuse audit evidence with workflow authority.
Audit	Infrastructure-like service candidate	Owns evidence of what happened; it should record, not decide business flow.
Reporting	Read-side service candidate	Owns projections and generated views; it is not source of truth.
Instrumentation	Conditional	High-volume facts and signal ownership matter, but timing and data volume require careful design.
DigitalTwin	Conditional	Model state and divergence are meaningful, but coupling to instrumentation and fleet must be tested.
ReinforcementLearning Advisory	Conditional/advisory	Recommendations are not commands; separation depends on safety of the advisory boundary.

This reading may change as the book develops. That is not weakness. It is the point of the process. Service boundaries should be argued into existence, not assumed.

Chapter checkpoint

You are ready to continue when you can answer these questions clearly.

- What does a bounded context protect that a microservice also must protect?
- What extra responsibilities does a microservice add beyond a bounded context?
- Why is a Web API controller not proof of a service boundary?
- Why should VerdictIssued and RootCauseVerdictIssued.v1 not be treated as the same event?
- When is the modular monolith still the better home for a context?
- Which NEXUS-1 sectors look like strong service candidates, and which are conditional?
- Why is a public contract more than a DTO?

HONEST BOUNDARY
This chapter has not yet assigned databases per service, API endpoints, broker topics, deployment units, or observability pipelines. It has only separated the language boundary from the runtime boundary. That separation is the necessary starting point for every later design decision.

Next chapter

Chapter 3 will examine the cost of distribution. The purpose is not to frighten the reader away from microservices, but to count the price honestly before deciding which NEXUS-1 boundaries are worth paying for.

A bounded context is the beginning of service reasoning, not the end of it.