

Spend Taxonomy: Structure for Transparency in Procurement Spend

The spend taxonomy is a uniform, hierarchical structure of all procurement volumes by commodity group, category and, where applicable, supplier or site. It creates a common language between procurement, controlling and business departments and forms the basis for reliable analyses, category strategies and cost reduction programmes. Without a clean taxonomy, spend management remains a collection of contradictory individual lists.

Key facts

- A spend taxonomy assigns all expenditure to multi-level categories, typically three to five levels deep.
- It is a prerequisite for reliable sourcing analyses, ABC analyses and category strategies.
- In practice, an existing standard classification is usually adapted to the company's reality rather than reinventing an entire classification system.
- Data quality determines the benefit: uncleaned free-text line items significantly distort the analysis.
- The taxonomy should be jointly maintained by procurement, controlling and business departments and updated regularly.
- It forms the basis for metrics such as Purchase Price Variance (PPV), savings validation and contract compliance.

Definition: spend taxonomy

A spend taxonomy is a company-wide binding category system with which all procurement and purchasing expenditure is systematically classified. It is hierarchically structured and maps all relevant procurement objects from broad commodity area to detail category.

Distinction from material groups and material classes

The spend taxonomy is not identical to the material master structure in the ERP. Item numbers and material classes serve operational disposition and bill-of-materials management, whereas the spend taxonomy is an analysis and steering layer. For example, a direct material item from production can be assigned to the same spend category as analogous services from the same technical family. What matters is a clean mapping via assignment rules, not a duplication of maintenance effort. The ideal scenario is a documented, versioned, unambiguous mapping between material classes and spend categories.

- Analysis layer overlaying the ERP structure
- one spend category can bundle several material classes
- document assignment rules instead of duplicating manually

Typical levels of a taxonomy

In practice, three to five hierarchy levels have proven effective. At the top are commodity groups such as

direct materials, indirect materials, capital goods or services. Below these come category families and concrete categories, often supplemented by a fourth level for subcategories. A hierarchy that is too deep increases maintenance effort, while one that is too shallow prevents targeted analyses. The right depth depends on company size, assortment breadth and the question of at which level category managers should take responsibility.

Structure and methods for creating a spend taxonomy

The build begins with a spend analysis based on raw data and ends with a binding governance process. Between these two steps, methodological principles determine the long-term viability of the system.

Standard classifications as a starting point

Rather than building a classification system from scratch, many procurement organisations draw on established standard classifications and adapt them. Best practice: keep the upper levels largely aligned with the standard and only design the lower levels in a company-specific way. This facilitates benchmark comparisons and exchange with suppliers who bring their own classifications. What is crucial is that deduplication and merge rules are defined so that duplicate categories with different names are avoided.

- align upper levels with the standard
- sharpen lower levels in a company-specific way
- define merge rules for duplicate categories

Data cleansing before classification

Before the actual assignment, the source data must be cleansed: supplier harmonisation, duplicate detection, enrichment of free-text line items and exclusion of payments without a purchase reference. Typically, 20 to 50 per cent of raw data comes from free text or consolidated line items and cannot be readily assigned. Proven aids include master-data-driven assignment rules, mapping tables and, increasingly, rule-based classification tools. A review by category experts nonetheless remains indispensable, particularly for mixed line items.

Governance and maintenance process

A taxonomy lives from continuous maintenance. It requires a clearly named owner, a change process for new categories and a versioned history. Otherwise, parallel shadow structures emerge within just a few years. An annual review cycle combined with an ongoing change request procedure has proven effective. New commodity groups, for example arising from future technologies or new manufacturing processes, should be able to be integrated quickly without breaking the history of analyses.

Application in operational and strategic procurement

The taxonomy is not an end in itself, but a working tool for daily disposition, strategic planning and supplier management.

Category strategies and ABC analysis

Based on the taxonomy, expenditure can be aggregated per category and prioritised using the ABC analysis: A categories receive their own category strategies and preferably A suppliers, while C categories are often standardised via catalogue solutions or framework agreements. Only the uniform category logic makes it visible which volumes can be bundled across sites and entities. This is the

prerequisite for synergy-driven sourcing projects and consistent framework agreements with suppliers.

Reconciliation with contract and master data

Combined with contract data, the taxonomy shows which volumes run under framework agreements and where contract compliance has gaps. Line items outside existing contracts can be addressed in a targeted manner, for example via a price deviation workflow for orders below agreed terms. The price variance on ordering can also be analysed per category and traced back to systematic causes. In the direct material environment, the taxonomy also supports processes such as APQP by assigning purchased parts per category to the responsible sourcing teams.

Supplier portfolios and compliance

In combination with supplier master data, category-supplier matrices emerge that make dependencies and single-source risks visible. Supplemented by sustainability data such as an ESG risk rating, compliance requirements can be prioritised per category: stricter requirements apply to risk categories, for example a supplier ESG improvement plan for conspicuous suppliers. This turns the taxonomy into a hub between sourcing, risk and sustainability management.

Key metrics and analyses based on the taxonomy

The measurable benefit of a spend taxonomy shows in analysis quality and speed of steering. Several metrics are suitable for assessing your own maturity.

Classification rate and data quality

The key maturity metric is the proportion of spend volume correctly classified. In mature organisations this typically exceeds 95 per cent; below that, every further insight suffers. In addition, the duplicate share in the supplier master and the proportion of unassignable line items should be measured regularly.

- classification rate typically > 95 % in mature organisations
- share of free-text line items as a cause indicator
- regular monitoring instead of one-off measurement

Per-category value contributions

Savings and price variances are only meaningfully demonstrated on a consistent category basis. Purchase Price Variance (PPV) can be analysed per category and compared against budget or market price references. Cost reduction targets in the annual plan are also anchored at taxonomy level and tracked in controlling. A uniform savings methodology is important, otherwise procurement and controlling argue over identical figures.

Link with supply chain analytics

In combination with supply chain analytics, the taxonomy yields proactive steering information: price index developments per category, supplier consolidation potential or the effects of lot size optimisation on spend volume. Logistics metrics such as on-time delivery in order processing can also be mirrored per category and linked to procurement decisions. In this way, a static structure becomes a dynamic steering instrument.

Risks, typical errors and controls

The most common errors arise not during creation but in sustained application. Those who know the following pitfalls can avoid them systematically.

Silo thinking and deviating classifications

If individual sites, entities or departments maintain their own structures, group-wide analyses lose their meaningfulness. Typical issues are deviating assignments of the same product per site or the double use of category names with different meanings. A binding central taxonomy with clear governance, supplemented by overlap-free assignment rules, provides the remedy. In group structures, the mapping of cross-entity purchasing is also relevant, for example when several units use contracts secured by a parent company guarantee.

Over-detailing and maintenance effort

An overly fine system creates high coordination and maintenance effort and is circumvented in practice by users. A simple rule has proven effective: only as much category detail as decisions build upon. Create subcategories only if a different sourcing behaviour, a different risk profile or different suppliers lie behind them. Everything else can be mapped via attributes such as material, standard or supply region.

Controls and auditability

Since the taxonomy is the starting point for reports and savings validation, it should be documented in an audit-proof manner. Assignment criteria, approval processes and change history must be recorded. For service categories with a Master Service Agreement (MSA), it must in particular be clarified whether effort-based services are consistently assigned to the correct category. Legal and contract-related analyses must always be reviewed separately on a case-by-case basis.

Current developments and outlook

Technology and sustainability requirements are noticeably changing the way spend taxonomies are handled. Three developments are shaping this.

Automated classification

Language models and rule-based tools today enable automatic assignment of free-text line items to spend categories with respectable hit rates in practice. This significantly reduces manual effort but does not replace expert review. A hybrid approach makes sense: automation for mass data, category experts for edge cases and new categories. Traceability of every automatic assignment remains essential.

Sustainability and ESG data in the taxonomy

Sustainability reporting and supply chain due diligence obligations require per-category visibility of risks. Spend categories are increasingly extended with ESG attributes such as risk country, industry risk or emissions relevance. This makes it possible to steer for which suppliers an ESG risk rating is obtained or a supplier ESG improvement plan is agreed. Tax and regulatory requirements must be reviewed on a case-by-case basis; the taxonomy provides the data foundation for this.

Integration into digital procurement platforms

Modern e-procurement and marketplace solutions use the taxonomy for suggestion logic: when

ordering, suitable categories, contracts and suppliers are automatically suggested. This increases contract compliance and reduces maverick buying. Automatic disposition of C parts is also controlled via category-specific rules, for example with safety stock and minimum stock per category. The taxonomy thus evolves from an analysis tool to an operational rulebook.

Typical metrics

Classification rate: > 95 %

Proportion of spend volume correctly assigned; below this threshold the meaningfulness of analyses drops significantly

Share of free-text line items: < 10 %

typical target value after cleansing; a high share is the main cause of misclassification

Covered contract spend: > 80 %

Proportion of volume under framework agreements; measurable per category via contract compliance

Categories with a strategy: 100 % of A categories

in conjunction with the ABC analysis; C categories often handled via catalogues or automatic disposition

Practical example

A mid-sized mechanical engineering company with several plants obtained contradictory volumes for purchased parts in a spend analysis because each plant maintained its own material groups. Savings metrics were hardly reliable, and contract utilisation could not be demonstrated. The company introduced a four-level spend taxonomy, based on a standard classification with company-specific subcategories.

- Raw data from all plants cleansed, suppliers harmonised, free-text line items assigned via mapping rules
- Classification rate increased from initially around 70 to over 96 per cent, owner and change process defined
- ABC analysis created per category; A categories received category strategies and bundled framework agreements
- PPV and contract compliance tracked per category in the monthly report

Conclusion

The spend taxonomy is the foundation of every data-driven procurement function: it unites material, goods and services expenditure under a common, hierarchical logic. Its benefit depends less on the perfect structure than on data quality, clear governance and consistent maintenance. Those who adapt standard classifications sensibly, consistently assign free text and integrate the taxonomy into processes such as ABC analysis, contract compliance and ESG assessment gain a powerful steering instrument. The effort typically pays off quickly through bundled volumes, reliable savings validation and faster sourcing decisions. Regular reviews keep the system viable even in the face of assortment or regulatory changes.

Frequently asked questions

What is a spend taxonomy?

A spend taxonomy is a company-wide, binding hierarchical category system for classifying all procurement expenditure. It typically structures volumes across three to five levels, from commodity groups down to detail categories. It forms the basis for spend analyses, category strategies and savings validation.

How can you recognise a good spend taxonomy?

A good taxonomy is free of overlaps, granular enough to support decisions, but not over-detailed. It is versioned, has a named owner and a change process. In practice, the classification rate of mature organisations typically exceeds 95 per cent.

How many levels should a spend taxonomy have?

Three to five hierarchy levels have proven effective in practice. What matters is that each category reflects a distinct sourcing behaviour or risk profile. Additional detail is better handled through attributes rather than additional levels.

What is the difference between a spend taxonomy and ERP material groups?

ERP material groups serve operational disposition and materials management, while the spend taxonomy serves analysis and strategic steering. A single spend category can bundle several material classes. The ideal scenario is a documented, unambiguous mapping between the two systems.

How do you build a spend taxonomy step by step?

First, spend data from all sources is cleansed and suppliers are harmonised. Then a standard classification is selected as a framework and tailored to the organisation, followed by classification and expert review. Finally, governance, an owner and a change process are made binding.

What role does the spend taxonomy play in savings validation?

Savings and price variances such as Purchase Price Variance (PPV) can only be reliably compared on a consistent category basis. The taxonomy ensures that before-and-after comparisons use the same category and the same scope. Without it, procurement and controlling typically argue over the underlying figures.