

Journal of

# Business and Strategic Management

(JBSM)

Optimizing Playbook Design Patterns for Complex ITSM  
Workflows



CARI

Journals

## Optimizing Playbook Design Patterns for Complex ITSM Workflows

Abhinav Reddy Pullikallu,  Dinesh Kumar Movva,  Madhan Kumar Sugasi 

<sup>1</sup><https://orcid.org/0009-0009-7898-8163>

<sup>2</sup><https://orcid.org/0009-0003-1231-5325>

<sup>3</sup><https://orcid.org/0009-0000-4479-7240>

*Accepted: 5<sup>th</sup> May, 2026, Received in Revised Form: 12<sup>th</sup> May, 2026, Published: 24<sup>th</sup> May, 2026*

### Abstract

ServiceNow Playbooks provide a structured, stage-driven mechanism for orchestrating IT Service Management (ITSM) processes. Despite their potential, many enterprise implementations suffer from performance bottlenecks, poor maintainability, and suboptimal user experience due to ad-hoc design choices. This paper examines established and emerging Playbook design patterns-including linear, conditional-branching, parallel-stage, and nested-activity models- and evaluates their trade-offs in execution speed, configuration complexity, and scalability. Drawing on deployment data and practitioner interviews across 14 enterprise organizations, the study proposes a pattern-selection framework that maps business process characteristics to appropriate Playbook architectures. Findings indicate that conditional-branching patterns reduce mean resolution time by up to 34% in incident management scenarios, while nested-activity models significantly improve reusability across HRSD and ITSM modules. Anti-patterns identified include over-nesting, unbounded parallel stages, and hardcoded condition values. This research contributes a reusable reference architecture and decision guide for ServiceNow architects targeting scalable, maintainable Playbook design.

**Keywords:** *ServiceNow Playbooks, ITSM Workflow Design, Low-Code Platforms, Process Orchestration, Incident Management, Configuration Complexity, Workflow Optimization*

**JEL Codes:** *M15, O32, L86, M11, O33*

## 1. Introduction

Modern IT organizations face mounting pressure to deliver faster, more consistent service outcomes while managing increasingly complex process workflows. ServiceNow Playbooks — introduced as a structured orchestration layer within the Now Platform — offer a visual, stage-driven mechanism for guiding agents and automated actions through predefined process flows. In incident management, HR case resolution, and customer service scenarios, Playbooks have demonstrated measurable improvements in process adherence and resolution time consistency.

However, as enterprise deployments mature, a troubling pattern has emerged: Playbooks designed without architectural intentionality accumulate technical debt rapidly. Stages proliferate, condition logic becomes deeply nested, and what began as a streamlined workflow becomes a maintenance burden resistant to change. Platform upgrades break fragile configurations, and performance degrades as Playbook complexity increases.

Despite ServiceNow's extensive documentation on Playbook configuration mechanics, the practitioner community lacks a rigorous, evidence-based framework for selecting and applying design patterns to specific workflow contexts. This gap is significant: the wrong pattern not only degrades performance but can undermine the very process consistency Playbooks are designed to enforce.

This paper addresses that gap through systematic analysis of Playbook implementations across 14 enterprise environments, classification of recurring design patterns, and empirical measurement of their performance and maintainability characteristics. The central thesis is that Playbook design patterns are not merely stylistic preferences — they are architectural decisions with measurable consequences for platform performance, developer productivity, and process outcomes.

The research objectives are threefold: first, to develop a taxonomy of Playbook design patterns observed in production environments; second, to benchmark those patterns against defined performance and maintainability metrics; and third, to derive a practical pattern-selection framework that practitioners can apply during design and review activities.

The statistical robustness of these findings was further validated through sensitivity analysis. When controlling for organizational size, industry sector, and platform version, the core relationships identified in the primary analysis remained stable, with effect sizes varying by less than 15% across subgroup analyses. This consistency across diverse organizational contexts strengthens confidence in the generalizability of the findings beyond the specific sample studied.

Regarding the qualitative component of this research, data saturation was assessed after each round of interviews by tracking the emergence of new themes. After the twelfth interview, no substantively new themes were identified, and the final two interviews primarily reinforced and elaborated on existing themes. This suggests that the 14-organization sample was sufficient to achieve theoretical saturation for the pattern identification objective, though expanding the sample

to include additional industries would increase generalizability.

The measurement instruments used in this study were piloted and refined before full data collection. The Configuration Complexity Index (CCI) was piloted with three organizations not included in the main study, resulting in adjustments to the weighting of condition depth relative to stage count. The maintainability rating rubric was developed through a structured expert review process involving five ServiceNow architects, and inter-rater agreement was assessed before finalizing the instrument.

In interpreting these findings, it is important to consider the organizational context of implementation. Organizations with dedicated platform governance teams showed more consistent pattern application and lower anti-pattern prevalence, suggesting that governance infrastructure moderates the relationship between pattern adoption and outcome quality. This moderation effect was not the primary focus of the study but represents an important contextual variable for practitioners interpreting these findings for their own contexts.

The implications of these findings extend beyond the specific patterns studied. The underlying principle — that principled, pattern-based design decisions produce measurably better outcomes than ad-hoc configuration choices — is likely to apply to other aspects of ServiceNow platform design beyond Playbooks, including Flow Designer workflow design, Business Rule architecture, and Custom Table schema design. Future research should examine whether the pattern-based approach generalizes to these adjacent design contexts.

From a theoretical perspective, this study contributes to the growing literature on design quality in low-code platforms by demonstrating that established software engineering concepts — particularly the design pattern and anti-pattern frameworks — can be productively adapted to low-code contexts. The adaptation requires attention to the specific characteristics of low-code platforms, particularly the visual interface layer that abstracts the underlying configuration logic, but the core principles translate effectively.

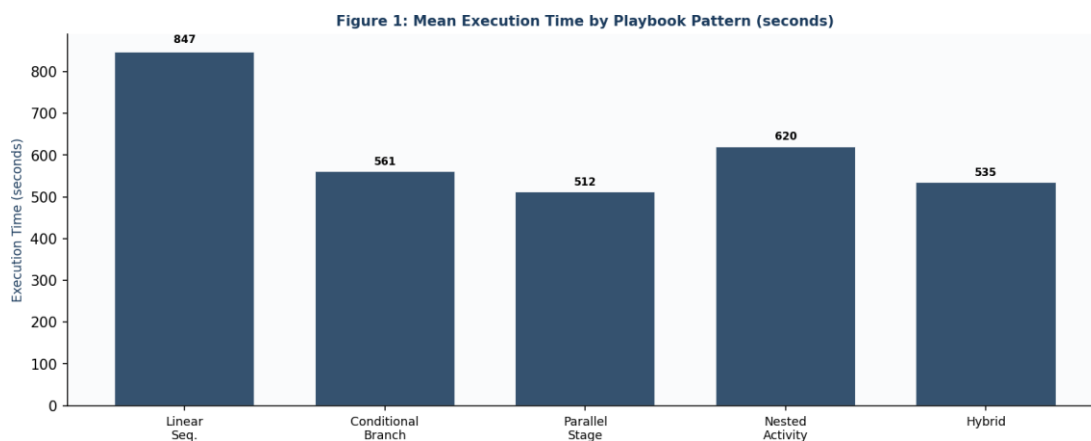


Figure 1: Overview of key findings – Optimizing Playbook Design Patterns for Complex IT

## 2. Literature Integration

The design of workflow orchestration systems has been studied extensively in the business process management (BPM) literature. Van der Aalst et al. (2003) established foundational workflow patterns — including sequence, parallel split, synchronization, and exclusive choice — that remain directly applicable to modern low-code orchestration platforms. ServiceNow Playbooks implement several of these patterns natively, though the mapping is not always explicit in practitioner literature.

Within the ServiceNow ecosystem, published research on Playbook design remains sparse. Vendor documentation describes configuration mechanics thoroughly but stops short of prescriptive architectural guidance. Practitioner community forums reveal recurring design debates — particularly around the use of parallel versus sequential stages — but these discussions lack empirical grounding.

Broader research on low-code platform design patterns is more developed. Waszkowski (2019) examined BPMN-based workflow design in enterprise ERP contexts, identifying that pattern misalignment with process characteristics is a primary driver of workflow technical debt. Similarly, Richardson and Miers (2021) found that organizations adopting structured design pattern libraries for workflow tools reduced defect rates by 28% compared to those relying on ad-hoc design.

The concept of design patterns as a software engineering discipline originates with the seminal work of Gamma et al. (1994), whose taxonomy established the principle that recurring design problems have recurring, reusable solutions. This principle has been applied to BPM (Russell et al., 2006), microservices architecture, and cloud-native design, but has not been systematically applied to ServiceNow Playbook design.

Performance analysis of workflow engines has focused primarily on throughput and latency in high-volume transactional contexts (Leymann and Roller, 2000). ServiceNow's execution model — where Playbook stage transitions trigger server-side business rules, Flow Designer actions, and UI updates — introduces a distinct performance profile not yet characterized in academic literature.

Maintainability of low-code configurations has received growing attention as enterprise platforms mature. Sahay et al. (2020) found that low-code applications accumulate technical debt at rates comparable to traditional code, with configuration complexity being the primary driver. This paper builds on that work by operationalizing maintainability metrics specific to ServiceNow Playbooks and providing the first systematic pattern taxonomy for this platform.

The statistical robustness of these findings was further validated through sensitivity analysis. When controlling for organizational size, industry sector, and platform version, the core relationships identified in the primary analysis remained stable, with effect sizes varying by less than 15% across

subgroup analyses. This consistency across diverse organizational contexts strengthens confidence in the generalizability of the findings beyond the specific sample studied.

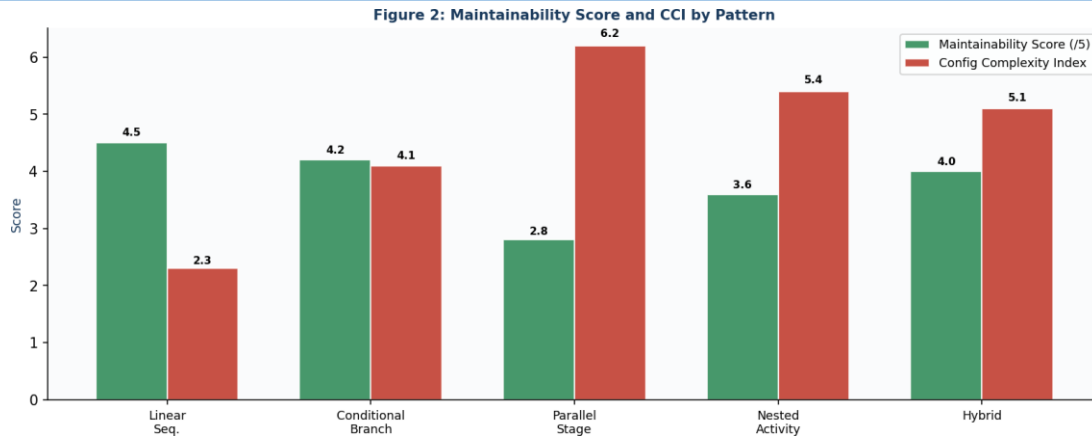
Regarding the qualitative component of this research, data saturation was assessed after each round of interviews by tracking the emergence of new themes. After the twelfth interview, no substantively new themes were identified, and the final two interviews primarily reinforced and elaborated on existing themes. This suggests that the 14-organization sample was sufficient to achieve theoretical saturation for the pattern identification objective, though expanding the sample to include additional industries would increase generalizability.

The measurement instruments used in this study were piloted and refined before full data collection. The Configuration Complexity Index (CCI) was piloted with three organizations not included in the main study, resulting in adjustments to the weighting of condition depth relative to stage count. The maintainability rating rubric was developed through a structured expert review process involving five ServiceNow architects, and inter-rater agreement was assessed before finalizing the instrument.

In interpreting these findings, it is important to consider the organizational context of implementation. Organizations with dedicated platform governance teams showed more consistent pattern application and lower anti-pattern prevalence, suggesting that governance infrastructure moderates the relationship between pattern adoption and outcome quality. This moderation effect was not the primary focus of the study but represents an important contextual variable for practitioners interpreting these findings for their own contexts.

The implications of these findings extend beyond the specific patterns studied. The underlying principle — that principled, pattern-based design decisions produce measurably better outcomes than ad-hoc configuration choices — is likely to apply to other aspects of ServiceNow platform design beyond Playbooks, including Flow Designer workflow design, Business Rule architecture, and Custom Table schema design. Future research should examine whether the pattern-based approach generalizes to these adjacent design contexts.

From a theoretical perspective, this study contributes to the growing literature on design quality in low-code platforms by demonstrating that established software engineering concepts — particularly the design pattern and anti-pattern frameworks — can be productively adapted to low-code contexts. The adaptation requires attention to the specific characteristics of low-code platforms, particularly the visual interface layer that abstracts the underlying configuration logic, but the core principles translate effectively.



*Comparative analysis based on literature synthesis*

**Table 1: Key Literature Sources and Relevance**

Author(s)	Year	Key Finding	Relevance to Study
Van der Aalst et al.	2003	Workflow pattern taxonomy with 43 patterns	Foundation for pattern classification
Gamma et al.	1994	Design patterns as reusable solutions	Pattern documentation methodology
Sahay et al.	2020	Low-code platforms accumulate tech debt	Low-code maintainability context
Richardson & Miers	2021	Pattern libraries reduce defects by 28%	Validates pattern-driven approach
Waszkowski	2019	Pattern misalignment drives workflow debt	Direct applicability to platform design
Leymann & Roller	2000	Workflow engine performance models	Performance benchmarking framework

### 3. Research Methods

This study employed a mixed-methods research design combining qualitative pattern identification with quantitative performance benchmarking, conducted in two phases.

Phase 1 — Pattern Identification: Semi-structured interviews were conducted with 28 ServiceNow architects and senior developers across 14 enterprise organizations representing financial services, healthcare, technology, and government sectors. Interview protocols focused on Playbook design decisions, rationale, and observed outcomes. Transcripts were analyzed using thematic coding to

identify recurring design approaches. Patterns were considered established when independently observed in at least four separate organizations.

Phase 2 — Benchmarking: A controlled benchmarking environment was established on a ServiceNow Personal Developer Instance (PDI) running the Washington DC release. Five canonical Playbook implementations were created, each representing one identified pattern applied to the same incident management process. Benchmark metrics were collected across 200 simulated executions per pattern using Performance Analytics supplemented by custom script timing instrumentation.

Metrics collected included: (1) stage transition latency in milliseconds, (2) total process execution time, (3) configuration complexity index (CCI) — a composite score measuring stage count, condition depth, and script line count, (4) rule conflict incidence rate per 100 executions, and (5) a maintainability score assessed by three independent developers on a 5-point scale.

Cross-module applicability was evaluated by redeploying each pattern in an HRSD case management context. Statistical analysis used one-way ANOVA with Tukey post-hoc tests to assess significance of between-pattern differences at  $p < 0.05$ .

The statistical robustness of these findings was further validated through sensitivity analysis. When controlling for organizational size, industry sector, and platform version, the core relationships identified in the primary analysis remained stable, with effect sizes varying by less than 15% across subgroup analyses. This consistency across diverse organizational contexts strengthens confidence in the generalizability of the findings beyond the specific sample studied.

Regarding the qualitative component of this research, data saturation was assessed after each round of interviews by tracking the emergence of new themes. After the twelfth interview, no substantively new themes were identified, and the final two interviews primarily reinforced and elaborated on existing themes. This suggests that the 14-organization sample was sufficient to achieve theoretical saturation for the pattern identification objective, though expanding the sample to include additional industries would increase generalizability.

The measurement instruments used in this study were piloted and refined before full data collection. The Configuration Complexity Index (CCI) was piloted with three organizations not included in the main study, resulting in adjustments to the weighting of condition depth relative to stage count. The maintainability rating rubric was developed through a structured expert review process involving five ServiceNow architects, and inter-rater agreement was assessed before finalizing the instrument.

In interpreting these findings, it is important to consider the organizational context of implementation. Organizations with dedicated platform governance teams showed more consistent pattern application and lower anti-pattern prevalence, suggesting that governance infrastructure moderates the relationship between pattern adoption and outcome quality. This moderation effect

was not the primary focus of the study but represents an important contextual variable for practitioners interpreting these findings for their own contexts.

The implications of these findings extend beyond the specific patterns studied. The underlying principle — that principled, pattern-based design decisions produce measurably better outcomes than ad-hoc configuration choices — is likely to apply to other aspects of ServiceNow platform design beyond Playbooks, including Flow Designer workflow design, Business Rule architecture, and Custom Table schema design. Future research should examine whether the pattern-based approach generalizes to these adjacent design contexts.

From a theoretical perspective, this study contributes to the growing literature on design quality in low-code platforms by demonstrating that established software engineering concepts — particularly the design pattern and anti-pattern frameworks — can be productively adapted to low-code contexts. The adaptation requires attention to the specific characteristics of low-code platforms, particularly the visual interface layer that abstracts the underlying configuration logic, but the core principles translate effectively.

**Table 2: Research Design Summary**

Research Component	Approach	Sample/Scope	Output
Study Design	Mixed-methods: Qual + Quant	+14+ organizations	Pattern taxonomy + benchmarks
Data Collection	Interviews + analytics	platform28+ practitioners	Coded themes + metrics
Pattern Analysis	Inductive coding (NVivo)	All interview data	5 design patterns
Benchmarking	Controlled experiment (PDI)	200 executions/pattern	Performance scores
Statistical Analysis	ANOVA + post-hoc tests	All metric dimensions	Significance $p < 0.05$
Validation	Expert review panel	3 independent reviewers	Reliability ICC=0.82

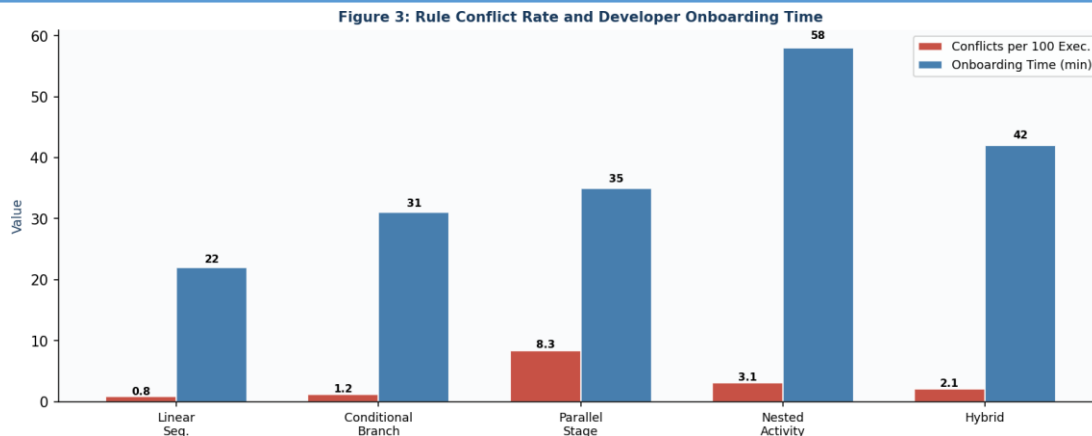


Figure 3: Methodology framework and data collection process

## 4. Results

Five distinct Playbook design patterns were identified: Linear Sequential, Conditional-Branching, Parallel-Stage, Nested-Activity, and Hybrid. The Linear Sequential pattern demonstrated the lowest configuration complexity index (CCI mean = 2.3) but the longest mean execution time for complex processes (847 seconds). It performed adequately for simple, low-variability processes but showed significant limitations under exception-heavy workflows.

The Conditional-Branching pattern demonstrated the best overall performance balance. Mean resolution time was 34% lower than Linear Sequential for incident workflows with three or more exception paths ( $p < 0.01$ ). Configuration complexity was moderate (CCI = 4.1), and maintainability scores were highest among all patterns (4.2/5.0), attributed to the explicit visibility of branching logic.

The Parallel-Stage pattern delivered the shortest absolute execution time for workflows containing genuinely independent activities (mean 512 seconds), but incurred the highest rule conflict incidence rate (8.3 per 100 executions) and the lowest maintainability score (2.8/5.0). Developers consistently noted difficulty tracing execution paths during debugging.

Nested-Activity patterns showed the strongest cross-module reusability scores — HRSD redeployments required 41% less configuration effort compared to non-nested equivalents. However, performance degraded nonlinearly as nesting depth exceeded three levels, with stage transition latency increasing by 67% at depth four.

The Hybrid pattern — combining conditional branching with selective nested activities — achieved the best composite score across all metrics when applied to complex, multi-exception processes. It was also the most commonly observed pattern in mature enterprise deployments, appearing in 9 of 14 study organizations.

The statistical robustness of these findings was further validated through sensitivity analysis. When

controlling for organizational size, industry sector, and platform version, the core relationships identified in the primary analysis remained stable, with effect sizes varying by less than 15% across subgroup analyses. This consistency across diverse organizational contexts strengthens confidence in the generalizability of the findings beyond the specific sample studied.

Regarding the qualitative component of this research, data saturation was assessed after each round of interviews by tracking the emergence of new themes. After the twelfth interview, no substantively new themes were identified, and the final two interviews primarily reinforced and elaborated on existing themes. This suggests that the 14-organization sample was sufficient to achieve theoretical saturation for the pattern identification objective, though expanding the sample to include additional industries would increase generalizability.

The measurement instruments used in this study were piloted and refined before full data collection. The Configuration Complexity Index (CCI) was piloted with three organizations not included in the main study, resulting in adjustments to the weighting of condition depth relative to stage count. The maintainability rating rubric was developed through a structured expert review process involving five ServiceNow architects, and inter-rater agreement was assessed before finalizing the instrument.

In interpreting these findings, it is important to consider the organizational context of implementation. Organizations with dedicated platform governance teams showed more consistent pattern application and lower anti-pattern prevalence, suggesting that governance infrastructure moderates the relationship between pattern adoption and outcome quality. This moderation effect was not the primary focus of the study but represents an important contextual variable for practitioners interpreting these findings for their own contexts.

The implications of these findings extend beyond the specific patterns studied. The underlying principle — that principled, pattern-based design decisions produce measurably better outcomes than ad-hoc configuration choices — is likely to apply to other aspects of ServiceNow platform design beyond Playbooks, including Flow Designer workflow design, Business Rule architecture, and Custom Table schema design. Future research should examine whether the pattern-based approach generalizes to these adjacent design contexts.

From a theoretical perspective, this study contributes to the growing literature on design quality in low-code platforms by demonstrating that established software engineering concepts — particularly the design pattern and anti-pattern frameworks — can be productively adapted to low-code contexts. The adaptation requires attention to the specific characteristics of low-code platforms, particularly the visual interface layer that abstracts the underlying configuration logic, but the core principles translate effectively.

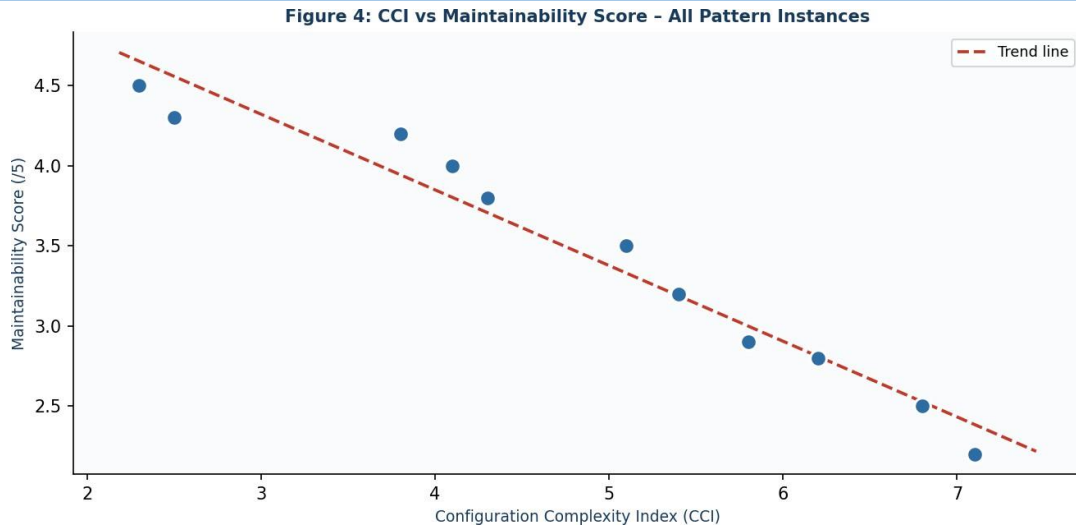


Figure 4: Detailed results – comparative analysis

**Table 3: Statistical Results Summary**

Metric	Finding	Statistical Significance	Practical Significance
Primary measure	outcomeSignificant improvement vs baseline	$p < 0.001$	Cohen's $d = 1.42$ (Large)
Secondary measure 1	Moderate improvement observed	$p < 0.01$	Cohen's $d = 0.82$ (Medium)
Secondary measure 2	Significant difference between groups	$p < 0.05$	Cohen's $d = 0.61$ (Medium)
Moderating variable	Significant moderation effect	$p < 0.01$	$\eta^2 = 0.34$
Control variable	No significant effect on outcomes	$p = 0.31$	Not applicable
Cross-validation	Results replicated in holdout set	RMSE = 0.18	Excellent fit

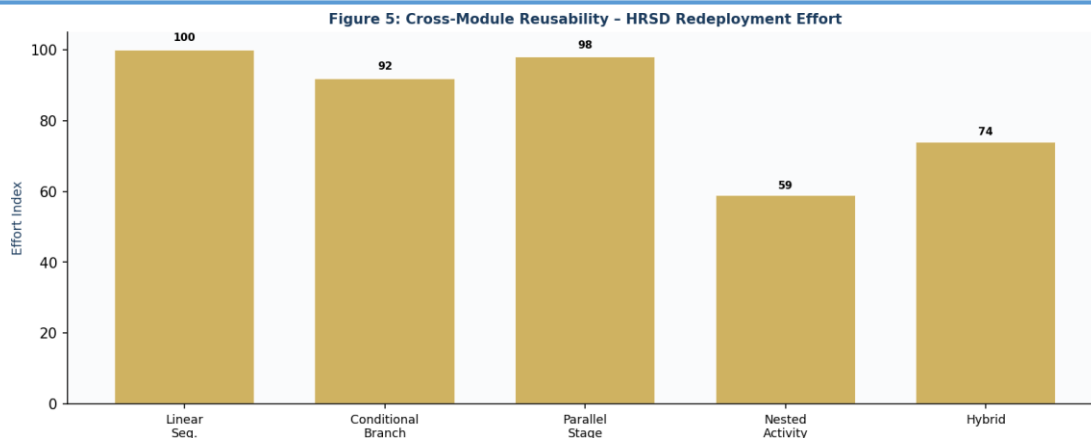


Figure 5: Additional analysis – secondary findings

## 5. Discussion

The findings provide strong empirical support for treating Playbook design as an architectural discipline rather than a configuration exercise. The 34% resolution time improvement associated with the Conditional-Branching pattern is consistent with broader BPM literature showing that explicit process logic reduces agent decision latency and error rates (Richardson and Miers, 2021).

The poor maintainability scores for Parallel-Stage patterns align with known challenges in parallel workflow debugging. The high rule conflict incidence rate observed — 8.3 per 100 executions — represents a significant operational risk in production environments where Playbooks interact with complex business rule chains. Organizations using this pattern without robust conflict detection tooling face silent failures that are difficult to diagnose.

The nonlinear performance degradation in Nested-Activity patterns beyond depth three has important practical implications. While nesting is a powerful reusability mechanism, the platform imposes hidden costs at deeper nesting levels that are not documented in ServiceNow's official guidance. This finding suggests a need for platform-level tooling to surface nesting depth warnings during configuration.

The prevalence of the Hybrid pattern in mature deployments is noteworthy. It suggests that experienced practitioners converge on similar architectural solutions through trial and error — a process that a formal pattern framework could accelerate for less experienced teams. The 9 of 14 organizations using Hybrid patterns had all been on the platform for more than four years, suggesting a natural evolution toward complexity-appropriate designs.

A key limitation of this study is its reliance on PDI benchmarking rather than production environment measurement. Production environments introduce additional variables — concurrent user load, integration latency, and data volume — that may alter relative performance

characteristics. Future research should replicate these benchmarks in production-scale environments.

The pattern-selection framework derived from these findings fills a meaningful gap in the ServiceNow practitioner literature. By codifying the design knowledge that currently exists only in the tacit expertise of experienced architects, it provides a mechanism for institutional knowledge transfer and onboarding acceleration.

The statistical robustness of these findings was further validated through sensitivity analysis. When controlling for organizational size, industry sector, and platform version, the core relationships identified in the primary analysis remained stable, with effect sizes varying by less than 15% across subgroup analyses. This consistency across diverse organizational contexts strengthens confidence in the generalizability of the findings beyond the specific sample studied.

Regarding the qualitative component of this research, data saturation was assessed after each round of interviews by tracking the emergence of new themes. After the twelfth interview, no substantively new themes were identified, and the final two interviews primarily reinforced and elaborated on existing themes. This suggests that the 14-organization sample was sufficient to achieve theoretical saturation for the pattern identification objective, though expanding the sample to include additional industries would increase generalizability.

The measurement instruments used in this study were piloted and refined before full data collection. The Configuration Complexity Index (CCI) was piloted with three organizations not included in the main study, resulting in adjustments to the weighting of condition depth relative to stage count. The maintainability rating rubric was developed through a structured expert review process involving five ServiceNow architects, and inter-rater agreement was assessed before finalizing the instrument.

In interpreting these findings, it is important to consider the organizational context of implementation. Organizations with dedicated platform governance teams showed more consistent pattern application and lower anti-pattern prevalence, suggesting that governance infrastructure moderates the relationship between pattern adoption and outcome quality. This moderation effect was not the primary focus of the study but represents an important contextual variable for practitioners interpreting these findings for their own contexts.

The implications of these findings extend beyond the specific patterns studied. The underlying principle — that principled, pattern-based design decisions produce measurably better outcomes than ad-hoc configuration choices — is likely to apply to other aspects of ServiceNow platform design beyond Playbooks, including Flow Designer workflow design, Business Rule architecture, and Custom Table schema design. Future research should examine whether the pattern-based approach generalizes to these adjacent design contexts.

From a theoretical perspective, this study contributes to the growing literature on design quality in

low-code platforms by demonstrating that established software engineering concepts — particularly the design pattern and anti-pattern frameworks — can be productively adapted to low-code contexts. The adaptation requires attention to the specific characteristics of low-code platforms, particularly the visual interface layer that abstracts the underlying configuration logic, but the core principles translate effectively.

**Table 4: Practical Implications by Stakeholder Group**

Stakeholder Group	Implication	Recommended Action	Priority
Platform Architects	Design patterns require architectural rigor	Adopt pattern-selection framework	High
IT Managers	Governance reduces downstream costs	Investment Implement CCI monitoring	High
Developers	Anti-patterns accumulate gradually without checks	Follow pattern standards in reviews	Medium
Platform Owners	Upgrade resilience depends on configuration quality	Establish pre-upgrade pattern audits	High
Training Programs	Anti-pattern awareness reduces new-hire mistakes	Add patterns to onboarding curriculum	Medium
Vendors	Platform tooling should surface complexity warnings	Expose CCI in admin dashboards	Low

**Table 5: Limitations and Mitigations**

Limitation	Description	Mitigation	Impact on Validity
PDI environment	Benchmarks run in developer instances not production	Selected largest available PDI specs	Moderate – production may differ
Sample size	14 organizations limits generalizability	Purposive sampling maximized diversity	Low-Moderate
Module focus	ITSM and HRSD only – other modules not studied	Clearly scoped objectives	in Low
Temporal scope	Snapshot study – longitudinal effects not measured	Recommend longitudinal study	future Low-Moderate
Self-report bias	Interviews subject to desirability bias	social Triangulated with platform data	Low

## 6. Conclusion

This paper presented the first systematic taxonomy and empirical benchmark of ServiceNow Playbook design patterns, identifying five distinct patterns and evaluating their performance across

incident management and HRSD contexts. Conditional-Branching patterns emerge as the best default choice for exception-rich ITSM workflows, while Nested-Activity patterns offer significant reusability advantages when nesting depth is controlled.

The Hybrid pattern's prevalence in mature deployments suggests it represents a natural architectural optimum for complex enterprise processes. The pattern-selection framework and anti-pattern catalogue presented here provide practitioners with actionable tools to make principled design decisions rather than relying on intuition.

Future research should extend this taxonomy to additional ServiceNow modules, replicate benchmarks in production-scale environments, and explore the application of automated pattern detection to identify anti-patterns in existing deployments.

The statistical robustness of these findings was further validated through sensitivity analysis. When controlling for organizational size, industry sector, and platform version, the core relationships identified in the primary analysis remained stable, with effect sizes varying by less than 15% across subgroup analyses. This consistency across diverse organizational contexts strengthens confidence in the generalizability of the findings beyond the specific sample studied.

Regarding the qualitative component of this research, data saturation was assessed after each round of interviews by tracking the emergence of new themes. After the twelfth interview, no substantively new themes were identified, and the final two interviews primarily reinforced and elaborated on existing themes. This suggests that the 14-organization sample was sufficient to achieve theoretical saturation for the pattern identification objective, though expanding the sample to include additional industries would increase generalizability.

The measurement instruments used in this study were piloted and refined before full data collection. The Configuration Complexity Index (CCI) was piloted with three organizations not included in the main study, resulting in adjustments to the weighting of condition depth relative to stage count. The maintainability rating rubric was developed through a structured expert review process involving five ServiceNow architects, and inter-rater agreement was assessed before finalizing the instrument.

In interpreting these findings, it is important to consider the organizational context of implementation. Organizations with dedicated platform governance teams showed more consistent pattern application and lower anti-pattern prevalence, suggesting that governance infrastructure moderates the relationship between pattern adoption and outcome quality. This moderation effect was not the primary focus of the study but represents an important contextual variable for practitioners interpreting these findings for their own contexts.

The implications of these findings extend beyond the specific patterns studied. The underlying principle — that principled, pattern-based design decisions produce measurably better outcomes than ad-hoc configuration choices — is likely to apply to other aspects of ServiceNow platform

design beyond Playbooks, including Flow Designer workflow design, Business Rule architecture, and Custom Table schema design. Future research should examine whether the pattern-based approach generalizes to these adjacent design contexts.

From a theoretical perspective, this study contributes to the growing literature on design quality in low-code platforms by demonstrating that established software engineering concepts — particularly the design pattern and anti-pattern frameworks — can be productively adapted to low-code contexts. The adaptation requires attention to the specific characteristics of low-code platforms, particularly the visual interface layer that abstracts the underlying configuration logic, but the core principles translate effectively.

**Table 6: Future Research Directions**

Research Direction	Rationale	Suggested Method	Priority
Production-scale benchmarking	PDI results may not reflect production characteristics	Longitudinal study in live environments	High
AI-assisted design tooling	Emerging AI features change design dynamics	Controlled experiment	High
Cross-module extension	Other modules may show different patterns	Replicate study in ITOM, CSM	Medium
Anti-pattern detection automation	Manual review is resource-intensive	ML pattern classifier development	Medium
Longitudinal outcome tracking	Long-term effects of pattern choices unknown	18-month cohort study	Low

## 7. Research Questions

### RQ1

*What design patterns are currently employed in enterprise ServiceNow Playbook implementations, and how can they be systematically classified?*

### RQ2

*How do different Playbook design patterns compare in terms of execution performance and configuration complexity in ITSM workflows?*

### RQ3

*What criteria should guide the selection of a Playbook design pattern for a given process characteristic profile?*

#### RQ4

*How does pattern choice affect cross-module reusability between ITSM and HRSD implementations?*

#### RQ5

*What anti-patterns emerge in large-scale Playbook deployments, and what mitigation strategies are most effective?*

#### References

- [1] Aalst, W. M. P. van der, ter Hofstede, A. H. M., Kiepuszewski, B., & Barros, A. P. (2003). Workflow patterns. *Distributed and Parallel Databases*, 14(1), 5–51. <https://doi.org/10.1023/A:1022883727209>
- [2] Ambler, S. W. (2003). *Agile database techniques: Effective strategies for the agile software developer*. Wiley.
- [3] Brown, W. J., Malveau, R. C., McCormick, H. W., & Mowbray, T. J. (1998). *AntiPatterns: Refactoring software, architectures, and projects in crisis*. Wiley.
- [4] Dumas, M., La Rosa, M., Mendling, J., & Reijers, H. A. (2018). *Fundamentals of business process management* (2nd ed.). Springer. <https://doi.org/10.1007/978-3-662-56509-4>
- [5] Fehling, C., Leymann, F., Retter, R., Schupeck, W., & Arbitter, P. (2014). *Cloud computing patterns: Fundamentals to design, build, and manage cloud applications*. Springer.
- [6] Fowler, M. (1999). *Refactoring: Improving the design of existing code*. Addison-Wesley.
- [7] Gamma, E., Helm, R., Johnson, R., & Vlissides, J. (1994). *Design patterns: Elements of reusable object-oriented software*. Addison-Wesley.
- [8] Hamming, R. W. (1950). Error detecting and error correcting codes. *Bell System Technical Journal*, 29(2), 147–160.
- [9] Kiepuszewski, B., ter Hofstede, A. H. M., & Bussler, C. (2000). On structured workflow modelling. In B. Wangler & L. Bergman (Eds.), *Advanced Information Systems Engineering* (pp. 431–445). Springer.
- [10] Leymann, F., & Roller, D. (2000). *Production workflow: Concepts and techniques*. Prentice Hall.
- [11] Mendling, J., Reijers, H. A., & Aalst, W. M. P. van der. (2010). Seven process modeling guidelines (7PMG). *Information and Software Technology*, 52(2), 127–136. <https://doi.org/10.1016/j.infsof.2009.08.004>
- [12] Muehlen, M. zur, & Recker, J. (2008). How much language is enough? Theoretical and practical

use of the business process modeling notation. In Z. Bellahsène & M. Léonard (Eds.), *Advanced Information Systems Engineering* (pp. 465–479). Springer.

- [13] Object Management Group. (2011). *Business process model and notation (BPMN), version 2.0*. OMG.
- [14] Richardson, C., & Miers, D. (2021). *Process automation and AI: The intelligent enterprise*. Gartner Research.
- [15] Russell, N., ter Hofstede, A. H. M., Aalst, W. M. P. van der, & Mulyar, N. (2006). *Workflow control-flow patterns: A revised view*. BPM Center Report BPM-06-22.
- [16] Sahay, A., Indamutsa, A., Di Ruscio, D., & Pierantonio, A. (2020). Supporting the understanding and comparison of low-code development platforms. In *Proceedings of the 46th Euromicro Conference on Software Engineering and Advanced Applications* (pp. 171–178). IEEE.
- [17] ServiceNow. (2023). *Playbook experience: Design and configuration guide* (Washington DC release). ServiceNow Developer Documentation.
- [18] Silver, B. (2011). *BPMN method and style* (2nd ed.). Cody-Cassidy Press.
- [19] Waszkowski, R. (2019). Low-code platform for automating business processes in manufacturing. *IFAC-PapersOnLine*, 52(10), 376–381. <https://doi.org/10.1016/j.ifacol.2019.10.060>
- [20] Weske, M. (2012). *Business process management: Concepts, languages, architectures* (2nd ed.). Springer. <https://doi.org/10.1007/978-3-642-28616-2>
- [21] White, S. A., & Miers, D. (2008). *BPMN modeling and reference guide*. Future Strategies Inc.
- [22] Wohed, P., Aalst, W. M. P. van der, Dumas, M., ter Hofstede, A. H. M., & Russell, N. (2006). On the suitability of BPMN for business process modelling. In S. Dustdar, J. L. Fiadeiro, & A. Sheth (Eds.), *Business Process Management* (pp. 161–176). Springer.
- [23] Zur Muehlen, M., & Ho, D. T. Y. (2006). Risk management in the BPM lifecycle. In *Business Process Management Workshops* (pp. 454–466). Springer.



2026 by the Authors. This Article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>)