HosPortal
CONNECTED HEALTHCARE SOLUTIONS

Maitland Emergency Department AI-Assisted Rostering

Whitepaper October 2025

Introduction

The implementation of the HosPortal rostering system at Maitland Emergency Department (ED) presented a complex and multifaceted scheduling challenge. The department's workforce is characterised by a diverse mix of staff, each with distinct rostering requirements and contractual conditions. The objective in selecting HosPortal was to leverage AI-driven automation to streamline scheduling while preserving fairness, flexibility, clinical safety, and operational efficiency.

This whitepaper outlines the key challenges encountered, the design and implementation approach adopted, the outcomes achieved, and the lessons learned in deploying an AI-assisted rostering solution within a high-acuity clinical environment.

Key challenges

Pattern Complexity

At the outset, the department maintained 28 legacy rostering patterns, which in combination produced over two million potential roster configurations. This level of combinatorial complexity:

- Reduced performance of the HosPortal AI and increased computation time
- Constrained the ability to iterate and refine rosters rapidly

Workforce Composition

The ED workforce comprised:

- Senior clinicians (FACEMs and other senior medical staff)
- Junior medical officers
- Trainees and registrars
- Interns and early-career doctors

Each cohort operated under distinct shift rules, supervision requirements, and training needs. These differences drove a high degree of heterogeneity in:

- Minimum and maximum hours
- Night shift frequency and recovery requirements
- Supervision ratios and skill mix across shifts

Collectively, this contributed to substantial rostering complexity and necessitated a nuanced, data-driven approach to configuration.

Roster Scope and Flexibility

The implementation encompassed:

- More than 25 distinct rosters
- 9 defined roles, many of which were interchangeable across rosters

This required HosPortal's AI to:

- Dynamically allocate staff across rosters and roles
- Respect clinical coverage requirements and skill mix
- Maintain fairness in distribution of less desirable shifts (e.g. nights, weekends)

A static, pattern-based model could not scale effectively to this level of complexity, prompting a shift in design philosophy.

Implementation approach

Key Design Principles

From the outset, the rostering framework was structured to maximise flexibility and support robust automation. This included:

- **Shift categorisation**: Shifts were grouped into day, evening, and night blocks, with clearly defined role and staffing requirements for each category.

- **Role-specific rosters**: Certain rosters were designed specifically for key functions (e.g. pod leaders), ensuring alignment with established staffing models and clinical governance requirements.
- **Configurability**: The design allowed for easy adjustment of parameters such as maximum consecutive shifts, preferred patterns, and limits on nights and weekends.

This upfront design work created a solid foundation for the AI, enabling it to operate within well-defined, clinically safe boundaries while still offering flexibility.

From patterns to rules

A critical design decision was to transition from a heavily pattern-based setup to a predominantly rule-based configuration. Instead of encoding desired sequences as static patterns, the team modelled outcomes as constraints and preferences.

In practice, the AI engine did not rely on the 28 legacy patterns as fixed inputs. Instead, it reconstructed those sequences dynamically through the underlying rule and constraint framework. This allowed the system to interpret each pattern's intent without being constrained by predefined templates. For Maitland, this led to the complete removal of fixed patterns. The solver generated compliant shift sequences directly from operational rules, improving performance, reducing configuration overhead, and supporting faster iteration across roster cycles.

This shift delivered several benefits:

- **Improved AI efficiency**: Reducing the reliance on fixed patterns significantly lowered the search space, enabling faster and more reliable solves.
- **Maintainability**: Rules could be adjusted incrementally (e.g. tightening weekend fairness or altering maximum consecutive nights) without redesigning entire patterns.
- **Scalability**: New roles, roster lines, or contractual conditions could be incorporated by updating rules, rather than proliferating additional patterns.

In effect, the system moved from "hard-coded" rostering behaviour to a flexible rule engine aligned with the department's evolving requirements.

What Worked Well

Well-Structured Design

The initial planning and roster architecture were deliberately crafted to support automation. Clear categorisation of shifts, roles, and constraints meant that the AI could operate with a high degree of structure and predictability while still accommodating real-world variability.

Strong Client Collaboration

There was continuous and constructive engagement between the Maitland ED leadership and the HosPortal implementation team. This included:

- Iterative reviews of interim AI designed rosters as patterns and rules were tested and refined, maximizing roster coverage and reducing vacancies, whilst maintaining rules around safe working hours and appropriate skills mix
- Joint refinement of rules and settings
- Alignment on timelines, priorities, and success criteria

This partnership approach ensured that the implementation remained clinically relevant and operationally practical.

Targeted AI Enhancements

During the project, a new constraint—referred to as a **shift isolation constraint**—was introduced. This was designed to:

- Encourage the solver to allocate staff to consecutive shifts where appropriate
- Better reflect the way Maitland ED conceptualises and manages its rosters
- Produce rosters that "feel intuitive" to end users, facilitating acceptance and trust in the AI-generated outcomes

This enhancement demonstrates how client feedback can directly inform the evolution of the AI model.

Positive Client Feedback

Early results from the automated build process were well received. As one client representative noted:

"The result from the autobuild was pretty good. I was happy with the sequences of shifts that it generated and the distribution of different types of shifts for each user."

This feedback validated the overall direction of the implementation and confirmed that the system was achieving a strong baseline configuration from which further refinements could be made.

Client Perspective

The client has approached the implementation with a forward-thinking and collaborative mindset, emphasising continuous optimisation rather than a one-off "set and forget" deployment:

"As we get more familiar with the autobuild, we will get better at tweaking the settings and rules to optimize the results."

This attitude aligns closely with HosPortal's partnership philosophy: building trusted, iterative relationships in which system configuration, AI behaviour, and operational processes co-evolve to support clinical and organisational goals.

Key Learnings and Considerations

The Maitland ED implementation has highlighted several key lessons relevant to AI-assisted rostering in complex clinical settings:

- **Early discovery is critical**
  A deep, early-stage understanding of client needs, constraints, and staffing models is essential to designing rosters that are both flexible and scalable, and that align with the AI foundation from the outset.
- **Rule-based configurations outperform pattern-heavy models**
  Transitioning from pattern-based to rule-based configurations significantly improves AI efficiency, maintainability, and adaptability, particularly where workforce and operational requirements are dynamic.
- **Feedback loops drive quality**
  Continuous feedback between the client and the HosPortal team is fundamental to refining automated builds. Each cycle provides data and insight that can be used to adjust constraints, rebalance fairness, and enhance solver performance.
- **Incremental refinement builds trust and autonomy**
  With each successful iteration, the system becomes better calibrated to the local context. Over time, the goal is for the client to be able to run and adjust autobuilds autonomously, using HosPortal support for complex changes or major model updates.

Conclusion

The Maitland Emergency Department implementation demonstrates the value of AI-assisted rostering in managing the complexity inherent in modern clinical environments. By combining:

- A robust, flexible roster design
- A shift from legacy patterns to intelligent rule-based constraints
- Purposeful AI enhancements informed by real-world usage
- A collaborative, iterative partnership with the client

Maitland ED and HosPortal have achieved meaningful technical and operational gains.

Ongoing work will focus on fine-tuning fairness around weekend and shift-type distribution, further optimising solver performance, and enabling the department to operate the system with increasing independence. The shared objective is a seamless, equitable, and sustainable rostering experience that supports both staff wellbeing and high-quality patient care.