WHITE PAPER

Extending Site Reliability Engineering to the Shop Floor

A Data-Centric Approach to Production Reliability Engineering

Chris Monchinski
Chief Technical Officer
InflexionPoint

inflexion point

Your goals. Our guidance. Shared success.

Contents

Summary 3

The Convergence of IT Reliability and OT Performance 4

From SRE to PRE: Redefining the “Service” 5

Production Reliability Budgets: The Core Innovation 6

Contractual Alignment: Operations vs Maintenance 8

Role of AI: From Prediction to Reliability Governance 9

The Role of Edge Orchestration with ZEDEDA 11

Reference Architecture: Data-Centric PRE 12

Closed-Loop Reliability: From Data to Action 13

Business Value 14

Conclusion 15

Summary

Site Reliability Engineering (SRE), pioneered by Google, has transformed the operation of digital systems by introducing a disciplined, metrics-driven approach to reliability using Service Level Objectives (SLOs), observability, and automation. These principles allow organizations to balance innovation velocity with system stability in highly complex, distributed environments.

These methodologies have remained largely confined to IT systems, however, while operational technology (OT) environments — despite increasing digitalization — continue to rely on traditional maintenance strategies and fragmented performance metrics. At the same time, the rise of AI-driven condition-based maintenance using platforms such as Sorba.ai has introduced powerful predictive capabilities, but without a unifying reliability governance model.

This paper introduces Production Reliability Engineering (PRE), a framework that extends SRE principles into OT environments, enabling unified reliability management across digital systems, physical assets, and production processes.

The Convergence of IT Reliability and OT Performance

Digital transformation is fundamentally reshaping industrial environments. Manufacturing systems are no longer isolated control systems—they are increasingly distributed, data-driven, and interconnected across enterprise and cloud domains. Technologies such as MQTT-based Unified Namespace (UNS), edge computing, and AI/ML have created a new class of cyber-physical systems that behave more like software platforms than traditional automation stacks.

Despite this transformation, a gap persists: IT systems have adopted SRE practices to ensure reliability at scale, while OT systems continue to rely on reactive maintenance, periodic preventive strategies, and siloed performance metrics. This disconnect leads to inefficiencies, misaligned decision-making, and an inability to fully leverage the value of industrial data.

To address this, organizations must move beyond maintenance-centric thinking and adopt a reliability engineering mindset for production systems. PRE provides this bridge, aligning IT and OT under a common, metrics-driven operational framework. Organizations can establish contractual reliability models, enforce production error budgets, and enable closed-loop optimization of manufacturing systems at scale by combining:

- Asset-level intelligence (Sorba.ai)
- System-level observability and AIOps (Ciroos.ai)
- Edge orchestration and lifecycle management (ZEDEDA)
- Data-centric architectures (UNS, historian, MES)

From SRE to PRE: Redefining the "Service"

At the core of SRE is the concept of a "service"—a unit of functionality delivered to users. In IT, this is typically an application or API. In manufacturing, however, the equivalent "service" is the production capability of a system—a line, cell, or unit operation delivering product at a defined rate and quality.

This shift in perspective is critical. It reframes reliability not as "machine uptime," but as the consistent delivery of production outcomes. For example, a machine may be running (available), but if it is producing off-spec product or operating below target throughput, it is effectively degrading the "service."

By redefining the service in this way, PRE enables the direct application of SRE constructs—such as SLOs and error budgets—to manufacturing systems, creating a unified language for reliability across IT and OT domain.

SRE vs PRE Conceptual Mapping

SRE Concept	Digital Context	PRE Equivalent	Physical Context
Service	Web/API service	Production system	Line / cell / unit operation
Availability	Uptime	Availability	OEE Availability
Latency	Response time	Cycle time	Takt / throughput
Error rate	Failed requests	Quality loss	Scrap / rework
Incident	Outage	Production event	Downtime / deviation
Observability	Logs/metrics/traces	Industrial data	Historian + MES + edge

Production Reliability Budgets: The Core Innovation

One of the most transformative ideas in SRE is the error budget, which defines the allowable level of system unreliability. This concept enables organizations to balance risk and innovation in a quantifiable way.

In manufacturing, this translates into the Production Reliability Budget, which defines acceptable limits for downtime, performance degradation and quality deviation. This approach introduces a fundamentally new operating model. Instead of striving for absolute uptime or rigid adherence to preventive schedules, organizations can make economic trade-offs between production output, equipment stress, and maintenance intervention.

For example, operations may choose to temporarily exceed optimal operating conditions to meet demand—knowingly consuming part of the reliability budget. Conversely, when the budget is exhausted, production must be constrained to allow recovery. This creates a dynamic, data-driven balance between performance and reliability.

Definition: Production Reliability Budget

- A Production Reliability Budget defines the allowable limits of:
- Downtime (availability loss)
- Performance degradation (throughput loss)
- Quality deviation (scrap/rework)

Example:

- Target OEE: 85%
- Allowable performance loss: 5%
- Allowable downtime: 6 hours/week

Operational Impact

Production Reliability Budgets transform decision-making from reactive to economically optimized reliability management.

Condition	Action
Within budget	Operations optimize throughput
Near limit	Increased maintenance readiness
Budget exceeded	Controlled slowdown + intervention

Contractual Alignment: Operations vs Maintenance

Historically, operations and maintenance have operated with competing objectives. Operations prioritizes throughput and schedule, while maintenance focuses on asset longevity and risk mitigation. This misalignment often results in reactive firefighting, inefficient maintenance cycles, and suboptimal asset utilization.

Service Level Agreements as a Framework

PRE introduces internal Service Level Agreements (SLAs) between these groups, formalizing shared accountability for production outcomes. By defining SLOs and associated reliability budgets, both functions operate against a common set of measurable objectives. This mirrors how SRE aligns development and operations teams in IT environments.

Example SLA		
Service: Filling Line 2	SLOs: Availability ≥ 92% MTBF ≥ 72 hours Scrap ≤ 2.5% Cycle time ≤ 1.2 sec/unit	
	Error Budget: 6 hours downtime/week 3% performance degradation	Outcome Shared accountability Quantified trade-offs Data-driven decision-making

This contractual model transforms organizational behavior. Maintenance becomes a strategic partner in achieving production targets, while operations gains visibility into the cost of pushing equipment beyond optimal conditions. The result is a collaborative, data-driven decision framework that mirrors the DevOps/SRE alignment seen in IT organizations.

Role of AI: From Prediction to Reliability Governance

Asset-Level Intelligence with Sorba.ai

Platforms such as **Sorba.ai** enable organizations to move beyond static maintenance schedules by applying machine learning to real-time and historical process data. These systems can detect early signs of equipment degradation, predict failures, and optimize process parameters.

However, while these models provide valuable predictions, they do not inherently determine when or how to act in the context of broader production objectives. Without integration into a reliability framework, predictive insights risk becoming isolated recommendations rather than actionable decisions.

Platforms such as Sorba.ai provide:

- Predictive failure detection (e.g., bearings, pumps, valves)
- Process optimization (e.g., DO control, energy minimization)
- Remaining useful life (RUL) estimation

System-Level Intelligence with Ciroos.ai

Platforms such as **Ciroos.ai** provide the system-level intelligence required to operationalize PRE. By correlating data across IT and OT systems, these platforms enable holistic observability and real-time assessment of system health relative to defined SLOs.

Ciroos.ai extends SRE practices into manufacturing by:

- Tracking production SLO compliance
- Correlating anomalies across systems
- Identifying root causes of performance degradation
- Triggering automated responses

This transforms raw data and predictions into reliability-aware operational decisions. Platforms such as Ciroos.ai extend SRE principles across systems:

- Cross-domain observability (IT + OT)
- Event correlation and root cause analysis
- SLO tracking and alerting
- Automated incident response

Combined Insight

The combination of Sorba.ai and Ciroos.ai enables a transition from predictive analytics to prescriptive reliability management. This is a critical step in achieving true PRE.

Layer	Function
Sorba.ai	Predicts asset and process behavior
Ciroos.ai	Governs system-level reliability

Together, they enable a Closed-loop Production Reliability Engineering. For example, consider this sequential flow scenario:

1	Sensor data → Sorba
2	Sorba predicts: Failure risk ↑ Efficiency ↓
3	Data flows into UNS / Historian
4	Ciroos correlates: "This will breach SLO in 6 hours"
5	System triggers: Maintenance work order Throughput adjustment
6	Error budget preserved

The Role of Edge Orchestration with ZEDEDA

A key challenge in scaling PRE across industrial environments is the deployment and lifecycle management of edge-based applications. Manufacturing environments are inherently distributed, with hundreds or thousands of edge nodes running critical workloads.

ZEDEDA addresses this challenge by providing a secure, scalable platform for managing edge infrastructure and applications.

ZEDEDA enables:

- Centralized deployment of containerized applications (e.g., Sorba models, MQTT brokers)
- Secure device onboarding and zero-trust architecture
- Remote updates and lifecycle management
- Resilient operation in disconnected environments

This capability is essential for PRE because it transforms the shop floor into a programmable, software-defined environment. Without such orchestration, PRE remains limited to pilot implementations. With ZEDEDA, it becomes an enterprise-scale operating model.

Reference Architecture: Data-Centric PRE

PRE is inherently aligned with modern data-centric architectures built around the Unified Namespace (UNS). This approach ensures that all systems—edge, OT, IT, and cloud—operate on a shared, real-time data fabric.
The architecture consists of layered capabilities:

- Edge Layer: Sensors, PLCs, and industrial devices
- Edge Orchestration Layer: ZEDEDA managing compute and applications
- Intelligence Layer: Sorba.ai models running at the edge
- Data Layer: MQTT/UNS, historian, MES
- AIOps Layer: Ciroos.ai for observability and SLO management
- Governance Layer: PRE policies, budgets, and SLAs

Layered Architecture

A layered approach ensures scalability, interoperability, and alignment with standards such as ISA-95. The full stack of tools enabling PRE could be as follows:

Layer	Role	Example
Governance	SLOs, error budgets	PRE framework
AIOps	System reliability mgmt	Ciroos.ai
Data Layer	UNS, historian, MES	Ignition, PI, HighByte
Intelligence	ML/AI models	Sorba.ai
Edge Orchestration	Deployment + lifecycle mgmt	ZEDEDA
Edge Compute	Runtime execution	IOT2050, IPCs
Assets	Physical systems	PLCs, sensors

Closed-Loop Reliability: From Data to Action

The true value of PRE lies in its ability to create a closed-loop system where data continuously drives action. In this model, predictions, observations, and governance rules are tightly integrated. For example:

- A Sorba model detects increasing vibration in a pump
- Data is published to the UNS and contextualized
- Ciroos identifies a projected SLO breach within hours
- A maintenance intervention is automatically triggered

This loop ensures that reliability is not just monitored, but actively managed and optimized in real time.

Business Value

PRE delivers value across multiple dimensions. Operationally, it improves OEE, reduces downtime, and optimizes resource utilization. Organizationally, it aligns teams and reduces inefficiencies caused by conflicting priorities. Strategically, it establishes a foundation for autonomous operations and AI-driven manufacturing. Most importantly, PRE transforms reliability from a cost center into a strategic lever for performance and competitiveness.

Operational Benefits	Organizational Benefits	Strategic Benefits
Increased OEE	Alignment between operations and maintenance	Foundation for autonomous operations
Reduced unplanned downtime	Data-driven governance	Scalable across sites
Optimized energy and resource usage	Reduced firefighting	Enables AI-driven decision-making at enterprise level

PRE as a Service: A New Operating Model

For operators looking to adopt PRE without significant disruption, PRE can be delivered as a structured offering: Production Reliability Engineering as a Service (PREaaS) that includes:

- SLO and reliability budget definition
- Data architecture and UNS implementation
- AI model deployment and tuning
- Edge orchestration with ZEDEDA
- AIOps integration with Ciroos
- Governance and reporting frameworks

This model enables organizations to adopt PRE incrementally while ensuring scalability and long-term sustainability.

Conclusion

SRE has revolutionized how complex, distributed digital systems are operated, by introducing:

- Measurable reliability
- Accountability through SLOs
- Automation driven by data

Extending SRE to the Shop Floor

Manufacturing now stands at a similar inflection point. PRE represents the next evolution of smart manufacturing, where reliability is engineered, orchestrated, and continuously optimized. By extending SRE principles into OT through Production Reliability Engineering, organizations can:

- Unify IT and OT reliability models
- Transform maintenance into a strategic function
- Enable closed-loop, AI-driven optimization of production systems

Implementing Production Reliability Engineering across your teams — and enabling them with platforms such as Sorba.ai, Ciroos.ai, and ZEDEDA — organizations can unify IT and OT, align operations and maintenance, and achieve closed-loop optimization of production systems.