

The EDC Success Blueprint

A step-by-step guide
to building your
Enterprise Data Cloud

Contents

Introduction	3
What is the Enterprise Data Cloud?	3
What's in this guide	3
How to use this guide	3
Chapter 1: Operational efficiency	4
Pillar overview: Increase efficiency	4
Maturity level introduction	4
Maturity path summary	5
Maturity path: Level 1 → Level 2	6
Maturity path: Level 2 → Level 3	11
Maturity path: Level 3 → Level 4	15
Maturity path: Level 4 → Level 5	19
Maturity path: Level 5	22
Chapter 2: Business continuity	23
Pillar overview: Strategic impact and metrics	23
Maturity level introduction	23
Maturity path summary	24
Maturity path: Level 1 → Level 2	25
Maturity path: Level 2 → Level 3	29
Maturity path: Level 3 → Level 4	32
Maturity path: Level 4 → Level 5	35
Maturity path: Level 5	38
Chapter 3: Application agility and automation	39
Pillar overview: Application agility and automation	39
Maturity level introduction	39
Maturity path summary	40
Maturity path: Level 1 → Level 2	41
Maturity path: Level 2 → Level 3	46
Maturity path: Level 3 → Level 4	51
Maturity path: Level 4 → Level 5	56
Maturity path: Level 5	61
Why Everpure	62
The platform built for tomorrow	62
What that means in practice	62
The results organizations are seeing	62

Introduction

Most infrastructure strategies fail the same way. They start with a vision of what good looks like, skip past the honest assessment of where things stand today, and hand teams a list of technology recommendations with no real thread connecting the two. The Enterprise Data Cloud (EDC) success blueprint was built to fix that.

This guide gives you a structured path from your current environment to a fully realized EDC. It is organized around the three outcomes that matter most to IT leaders: reducing risk, increasing efficiency, and scaling with agility. Each chapter covers one pillar in depth, with prescriptive maturity steps, specific technology and process actions, and clear measures of progress.

You won't find generic best practices in this guide. Every recommendation is tied to where you are now and sequenced to move you forward without rebuilding everything at once.

What is the Enterprise Data Cloud?

The EDC is not a product. It is an architecture, a model built on the Everpure™ Platform for managing data as a unified, governed asset across on-premises, hybrid, and cloud environments. Where legacy infrastructure manages arrays, the EDC manages data. It understands how workloads consume data, automatically applies global policies, and governs placement, protection, and lifecycle decisions without manual intervention.

Enterprise IT spent decades building infrastructure around applications. Each system brought its own data model, governance rules, and storage requirements. The result is silos at every layer: application, semantic, infrastructure, and governance. AI exposes all of them at once.

The EDC closes that gap. It delivers a unified, intelligent data management layer across your entire environment so data is ready, governed, and available when the business needs it.

What's in this guide

This blueprint includes three chapters, each mapped to a specific business outcome:

- **Increase efficiency:** Reduce operating costs through automation, consolidation, and lifecycle management.
- **Business continuity:** Align disaster recovery, failover automation, and resilience governance to the risk exposure reduction outcome.
- **Scale with agility:** Automate application deployment and workload intelligence for on-demand business responsiveness.

Each chapter maps five maturity levels, from manual and reactive at Level 1 to fully autonomous and outcome-aligned at Level 5. The blueprint explains what changes at each level across technology, process, and governance, and which Everpure capabilities move you there.

How to use this guide

Start with the chapter that matches your highest-priority outcome. If you've already completed the EDC maturity assessment, that result tells you exactly which chapter to open first and which maturity level to target.

If you haven't taken the [assessment](#) yet, it's a three-question exercise that maps your current environment to one of the three outcome areas. It takes about two minutes to complete and gives you a personalized recommendation for where to start.

Work through the maturity steps in sequence. Each level builds on the last. Skipping ahead is possible, but the prerequisite actions at each stage exist for a reason: they're the ones that tend to surface as technical debt later if skipped.

By the time you've worked through the blueprint, you'll have a precise map of where your environment stands and a clear, sequenced plan to get it where it needs to be.

Chapter 1: Operational efficiency

True organizational efficiency is the engine that converts digital, financial, and operational capital into a formidable market advantage. It is the pulse of an enterprise, measuring how quickly and consistently a team can turn complex transformations into tangible value, slashing the time between an initial investment and the moment it begins to benefit customers and stakeholders alike.

Innovation alone is statistically a hollow victory, whereas strategic dominance requires the raw speed of execution and a business model that is perfectly in sync with its vision. When an organization falters here, it does more than just miss a deadline; it delays market responses, creates pockets of stranded investment, and erodes the vital trust that investors place in its growth potential.

In a high-performing culture, planning, budgeting, and execution move in a unified rhythm. We prioritize initiatives not just by their return on investment (ROI), but by their speed to impact. This ensures that every new automation effort or service line is launched with deep cross-functional accountability, where all team members understand the outcome they are working to achieve.

In today's business world, the cost of inaction is steep. It manifests as failed transformations, missed pivots, and a market perception that the organization simply isn't ready for the future. Beyond the balance sheet, these failures ripple through your brand and operations. Financially, it means forfeiting early-mover revenue and watching the potential of AI and automation slip away as competitors move to monetize these strategic shifts first.

Ultimately, the greatest risk is watching others capitalize on shifting cost structures and market dynamics while your own value realization lags behind. When roadmap ambitions aren't met with real-world results, even the most loyal analysts and investors begin to lose confidence in leadership's ability to deliver.

Pillar overview: Increase efficiency

- **Impact:** Reduce operating costs through automation and consolidation.
- **Target metric:** Reduce operating costs by 74% and achieve 85% fleet utilization (driving optimal unit economics).
- **Library focus:** Automation, lifecycle management, elasticity, scalability, and operational productivity

Maturity level introduction

The gap between Level 1 and Level 5 in this pillar is the gap between a team that spends most of its time on tickets and a platform that resolves most issues before a ticket is ever opened. That's not an exaggeration. At Level 1, every provisioning request, performance issue, and capacity question is a manual task. At Level 5, the system handles the routine work autonomously, and the team focuses on exceptions and business outcomes.

The following five maturity levels describe that progression. Each one is a distinct operational state with its own characteristics, not just a degree of the same capability. The difference between Level 2 and Level 3 is the difference between reducing manual work and actually automating it. The difference between Level 4 and Level 5 is the difference between predicting problems and eliminating them before they register.

Most organizations are at Level 2 or early Level 3. The tooling exists to go further. What tends to hold teams back is the process and governance work that needs to happen alongside the technology. This blueprint is designed to unblock exactly that.

Maturity levels

Level	Maturity description
ML-1	- Storage tasks are manually performed across siloed systems. - Troubleshooting and provisioning are ticket-based and time-consuming.
ML-2	- Scripts and admin tools reduce workload but require platform expertise. - Performance issues are addressed post-failure.
ML-3	- Repeatable tasks are automated via playbooks. - Monitoring tools surface trends and anomalies for intervention.
ML-4	- AI-based recommendations improve performance and reduce overprovisioning. - Storage tiering and provisioning are policy-guided.
ML-5	- Ops run autonomously across storage domains. - CIOs monitor service level objective (SLO) adherence, auto-remediation frequency, and backlog elimination rates.

Maturity path summary

The following table maps all four maturity transitions in a single view. Each row shows what needs to change across technology, process, and governance, along with the Everpure capabilities that enable movement from one level to the next. It's a fast way to orient yourself before stepping into the detailed guidance for each transition.

There are three things worth noting before you use it. First, the columns don't always move together. Some organizations have strong monitoring but no process accountability around what the monitoring surfaces. Others have governance structures in place but haven't connected them to the right tooling yet. The table helps you see the full picture, not just the technology gap. Second, the sequencing is intentional. The automation at Level 3 only works cleanly if the ownership and runbooks from Level 2 are already in place. Skipping steps is possible, but it tends to create technical debt that shows up later. Third, you don't have to start at Level 1. Enter at the level that reflects where your organization actually is today.

Paths

Maturity transition	Technology enhancements	Process maturity	Governance	Enabling features
Level 1 → 2	Introduce monitoring for performance baselines and implement manual alerting thresholds.	Define operational runbooks for capacity planning and maintenance; manually track usage anomalies.	Assign system ownership; introduce periodic performance reviews across key platforms.	Pure1® basic monitoring, manual service level agreement (SLA) thresholds, and Everpure Fusion™ log review templates
Level 2 → 3	Deploy integrated observability tools and service level metrics; alert on deviations from baselines.	Automate ticket routing and impact diagnostics; enforce regular review of performance and uptime.	Centralize ops coordination across teams; define escalation workflows based on SLA tiers.	Pure1 insights engine, Everpure Fusion capacity planners, and alert correlation tooling
Level 3 → 4	Enable predictive analytics for anomaly detection and automated resource tuning.	Shift operations toward exception-based handling; implement continuous improvement loops.	Operational key performance indicators (KPIs) are shared across engineering, site reliability engineering (SRE), and infrastructure teams, with executive oversight of trends.	Pure1 predictive engine, Everpure Fusion trend modeling, and exception-handling triggers
Level 4 → 5	Adopt AI-assisted auto-remediation and real-time workload rebalancing infrastructure.	Run autonomous remediation for common issues; apply intent-based operations across platforms.	Efficiency and availability are embedded into objectives and key results (OKRs) across business units; governance includes real-time ops.	Everpure Fusion AI remediation modules, Pure1 intent-based recommendations, and real-time ops scoring

Maturity path: Level 1 → Level 2

In ML-1, the storage tasks are manually performed across siloed systems, while troubleshooting and provisioning are ticket-based and time-consuming.

Category	Requirement
Technology enhancements	Introduce monitoring for performance baselines and implement manual alerting thresholds.
Process maturity	Define operational runbooks for capacity planning and maintenance; manually track usage anomalies.
Governance and measurement	Assign system ownership; introduce periodic performance reviews across key platforms.
Benefits and metrics	Support tickets are driven by reactive issues, with limited insights into root causes or patterns.
Everpure features	Pure1 basic monitoring, manual SLA thresholds, and Everpure Fusion log review templates

Technology enhancements

Requirement: Introduce monitoring for performance baselines and implement manual alerting thresholds.

Recommended actions:

- Implement centralized cost visibility tools and basic tagging across all environments.
- Deploy performance and capacity monitoring agents on all storage arrays and related infrastructure.
- Implement Pure1 API-first telemetry integration for 100% data capture, ensuring a unified data model (for example, JSON structure) for future aggregation and pipeline ingestion.
- Deploy an analytical reporting solution such as Power BI or Grafana, and set up basic metrics reporting.
- Establish performance baselines for key platforms to enable anomaly detection and overall metrics.

Who to involve:

- **Storage administrator:** Deploys monitoring agents, configures alerting thresholds, and ensures 100% environment telemetry reporting to Pure1, security information and event management (SIEM) systems, or other collection frameworks
- **Infrastructure architect:** Defines performance baselines for key platforms and establishes the structure for the basic tagging policy to ensure alignment with wider infrastructure strategy
- **Financial analyst:** Collaborates with the operations team to implement centralized cost visibility tools and interpret initial performance/cost dashboards

Success indicators:

- 100% of the storage environment is inventoried and actively reports telemetry data to Pure1, SIEM, site reliability, or other collection frameworks.
- Initial critical performance thresholds are defined and configured for manual alerting.
- All key workloads are tagged for basic spend categorization (for example, env=prod, billingID=FIN).

Everpure enablers:

- Pure1 basic array and fleet monitoring
- Pure1 metrics and cost-based storage dashboards
- Pure1 workload planner for capacity and cost-based storage dashboards

Process maturity

Requirement: Define operational runbooks for capacity planning and maintenance; manually track usage anomalies.

Recommended actions:

- Formalize a monthly reporting cadence for cost trends and capacity utilization.
- Develop a standardized, hierarchical tagging policy (for example, BU, Env, App) for all new storage provisioning and associated infrastructure to ensure granular cost attribution.
- Codify runbooks in a version-controlled repository (for example, Git/Confluence) to establish a single source of truth for operations, enabling future infrastructure-as-code conversion.
- Create a root cause analysis process that tightly integrates with a configuration management database (CMDB), alerting, and ticketing solutions.

Who to involve:

- **Storage administrator:** Implements the standardized tagging policy and documents runbooks and standard operating procedures for routine capacity and maintenance operations
- **IT service manager/IT operations lead:** Drives the formalization of the monthly reporting cadence and defines the structure of the root cause analysis process for service continuity

Success indicators:

- Monthly cost trend reports are consumed by operations, application owners, and finance teams.
- New workload provisioning is compliant with the standard tagging policy (>95% allocation coverage).
- Usage anomalies are identified and manually tracked for key workloads.

Everpure enablers:

- Everpure Fusion workload labels and Purity tags
- Everpure Fusion and Pure1 capacity planner AI to support manual capacity plan input
- Pure1 tagging API to propagate tags to volumes programmatically

Governance and measurement

Requirement: Assign system ownership; introduce periodic performance reviews across key platforms.

Recommended actions:

- Establish clear ownership and accountability for storage consumption and spend by environment.
- Implement a “showback” reporting model to present allocated spend by team and environment.
- Define platform-specific SLOs and metrics for periodic performance and uptime reviews.

Who to involve:

- **Head of Infrastructure:** Drives the centralization of operational coordination and formally defines escalation workflows based on SLA tiers
- **Application owner/line-of-business (LoB) lead:** Assigns ownership and accountability for storage consumption and spend by environment, and utilizes showback reports to manage application budgets
- **Financial analyst:** Defines the structure of the showback reporting model and interprets initial performance and cost dashboards to align IT spend with financial targets
- **Head of Infrastructure:** Provides executive validation of the showback model’s accuracy and formally signs off on escalation workflows based on SLA tiers
- **Financial analyst:** Leads the integration of cost forecasting with application SLAs to ensure that spending aligns with defined service tiers and financial targets

Success indicators:

- Budget accountability is assigned by the environment and tied to the platform team.
- Initial showback dashboards are published to build trust in cost data.
- Service level metrics (for example, latency, throughput) are gathered and reviewed periodically.

Everpure enablers:

- Pure1 dashboard prioritizer for cost dashboards
- Evergreen//One™ adjustable service tiers for cost model understanding
- Everpure Fusion financial correlation viewer to map usage to cost

Benefits and metrics

Requirement: Support tickets are driven by reactive issues, with limited insights into root causes or patterns.

Recommended actions:

- Establish baseline metrics for time to resolve (TTR) and number of reactive support tickets.
- Begin tracking capacity utilization metrics to identify initial consolidation opportunities.
- Track the total number of manual operational tasks per month to quantify efficiency debt.

Who to involve:

- **Site reliability engineer:** Defines the technical metrics baseline for TTR and reactive ticket volume tracking, ensuring data collection aligns with reliability standards
- **Storage administrator:** Compiles the raw data on TTR, ticket volume, and capacity utilization and uses the data to identify initial consolidation opportunities
- **IT service manager/IT operations lead:** Leverages TTR and manual task data to measure efficiency debt and champion process improvements, ensuring metrics support service continuity goals

Success indicators:

- Initial TTR and reactive ticket volume baselines are recorded.
- Consolidation projects are identified using capacity utilization data.
- Data is available to forecast variances in capacity and performance.

Everpure enablers:

- Pure1 volume metrics for utilization tracking
- Everpure Fusion capacity planner to model growth and consolidation
- Everpure Fusion lifecycle risk modeler to track manual process risk

Everpure features

Requirement: Pure1 basic monitoring, Everpure Fusion log review templates, and API-first telemetry integration

Recommended actions:

- Leverage the Pure1 tagging API to propagate tags to volumes for initial cost attribution.
- Configure Pure1 volume metrics to capture baseline utilization and performance data.
- Use Everpure Fusion log review templates to establish manual tracking of operational anomalies.

Who to involve:

- **Storage administrator:** Configures and monitors Pure1 volume metrics and ensures successful tag propagation via the Pure1 tagging API for initial cost attribution and data collection
- **Financial analyst:** Confirms that tag propagation via the Pure1 tagging API aligns with the required financial categorization for initial cost attribution reporting

Success indicators:

- Pure1 dashboards reflect current volume metrics and cost-based views.
- Tags propagated via the Pure1 tagging API are visible and used for basic reporting.
- Log review templates are used to document usage anomalies for manual intervention.

Everpure enablers:

- Pure1 basic monitoring
- Pure1 volume metrics
- Pure1 tagging API
- Everpure Fusion capacity planner
- Everpure Fusion lifecycle risk modeler

Maternity path: Level 2 → Level 3

Transitioning to ML-3 includes the use of scripts and admin tools to reduce workload, but it requires platform expertise; performance issues are addressed post-failure.

Category	Requirement
Technology enhancements	Deploy integrated observability tools and service level metrics; alert on deviations from baselines.
Process maturity	Automate ticket routing and impact diagnostics; enforce regular review of performance and uptime.
Governance and measurement	Centralize operations coordination across teams; define escalation workflows based on SLA tiers.
Benefits and metrics	TTR and capacity saturation metrics are actively tracked; baseline performance is maintained.
Everpure features	Pure1 insights engine, Everpure Fusion capacity planners, and alert correlation tooling

Technology enhancements

Requirement: Deploy integrated observability tools and service level metrics; alert on deviations from baselines.

Recommended actions:

- Deploy a lightweight event correlation engine to process alerts from Pure1/SIEM, automating incident suppression and prioritizing actionable deviation signals over isolated events.
- Integrate storage telemetry data with broader observability tools for full service level visibility.
- Automate ticket creation and initial routing based on critical system alerts and ensure synchronization with the CMDB.

Who to involve:

- **Storage engineer/administrator:** Integrates storage telemetry data with broader observability tools and implements automated monitoring and alerts based on deviations from baselines
- **Site reliability engineer:** Focuses on performance and availability, developing scripts for automated ticket routing, and initial impact diagnostics to reduce mean time to identify (MTTI)
- **Head of Infrastructure:** Centralizes operations coordination across teams and formally defines SLAs to be tracked
- **IT service manager/CMDB owner:** Ensures the CMDB is accurately synchronized with automated ticket creation and routing to maintain a reliable configuration and asset repository

Success indicators:

- Monitoring is automated to flag performance degradation before failures occur.
- Time spent on manually tracking and routing tickets is measurably reduced.
- Performance and capacity metrics are actively tracked and reported against baseline.

Everpure enablers:

- Pure1 insights engine
- Pure1 alert rules and alert correlation tooling
- Everpure Fusion capacity planners for active monitoring against planned growth

Process maturity

Requirement: Automate ticket routing and impact diagnostics; enforce regular review of performance and uptime.

Recommended actions:

- Develop automated diagnostics scripts for common performance issues to reduce MTI.
- Standardize incident response workflows and integrate them with centralized operations coordination.
- Enforce a governance process for regular review of performance and uptime metrics across all platforms.

Who to involve:

- **Storage engineer/administrator:** Develops operations, deploys automated diagnostic scripts for common storage issues to reduce MTI, and standardizes incident response
- **IT service manager:** Standardizes and integrates incident response workflows with centralized operational coordination to ensure predictable issue resolution

Success indicators:

- MTI and TTR are reduced by automated diagnostics.
- Ticket routing and initial impact assessment are automated.
- Performance reviews lead to documented capacity adjustments and optimization actions.

Everpure enablers:

- Everpure Fusion capacity planner to monitor and optimize capacity
- Pure1 ticket routing model
- Everpure Fusion SLA severity index for prioritizing issues

Governance and measurement

Requirement: Centralize operations coordination across teams; define escalation workflows based on SLA tiers.

Recommended actions:

- Centralize operational coordination and reporting across siloed infrastructure and application teams.
- Formally define and implement SLAs for different application tiers.
- Implement escalation workflows that automatically trigger based on SLA adherence deviation.
- Align cost forecasting capabilities with application-level SLAs to ensure spend is justified by service tier.

Who to involve:

- **Enterprise application architect:** Formally defines the SLAs for different application tiers to ensure infrastructure services meet business expectations
- **Head of Infrastructure:** Provides executive sign-off on the formal SLAs and ensures the centralization of operational coordination is effective across siloed teams
- **Financial analyst:** Collaborates on formalizing the monthly reporting cadence for cost trends and validates the hierarchical tagging policy for accurate cost attribution and reporting
- **Application owner/LoB lead:** Validates that the tagging policy accurately reflects application ownership and consumption for proper cost attribution

Success indicators:

- Cross-team operational coordination reduces time to delivery for new provisioning requests.
- Escalations are automatically routed according to predefined SLA tiers.
- Cost forecasting is tied to workload profiles and aligns spend with application-level SLAs (ML-3 state).

Everpure enablers:

- Everpure Fusion SLA severity index
- Pure1 metadata scanning for centralized governance data
- Evergreen//One dashboards to track consumption against SLA/spend

Benefits and metrics

Requirement: TTR and capacity saturation metrics are actively tracked; baseline performance is maintained.

Recommended actions:

- Begin tracking MTTI to measure diagnostic automation effectiveness.
- Correlate storage operational efficiency (for example, automation rate) with application-level SLAs.
- Track the automation coverage percentage as an operational KPI—automated tasks ÷ total tasks—to set a measurable target for the transition to exception-based operations.

Who to involve:

- **Automation engineer/DevOps:** Provides data on the automation rate and calculates the cost per manual operation to quantify the ROI for automated processes
- **Financial analyst:** Correlates storage operational efficiency metrics (such as automation rate) with application-level SLAs to provide financial justification for ongoing efficiency investments

Success indicators:

- MTTI is reduced quarter over quarter through improved diagnostics.
- Cost per operation is quantified, establishing a target for future automation ROI.
- Utilization rates increase toward the target of 85% across the storage fleet.

Everpure enablers:

- Pure1 volume metrics for capacity and performance tracking
- API-first automation for measuring automation rate
- Everpure Fusion capacity planner to measure saturation against forecasts

Everpure features

Requirement: Pure1 insights engine, Everpure Fusion capacity planners, and alert correlation tooling

Recommended actions:

- Deploy Pure1 alert rules and correlation tooling to automatically flag deviations from baseline performance.
- Implement the Pure1 ticket routing model to automate initial ticket creation and prioritization based on severity.
- Utilize Everpure Fusion and Pure1 capacity planners for active monitoring against planned growth and consolidation forecasts.

Who to involve:

- **Storage engineer/administrator:** Deploys Pure1 alert rules, implements the ticket routing model, and utilizes Everpure Fusion and Pure1 capacity planners to proactively manage growth and prevent performance degradation

Success indicators:

- Critical performance deviations are automatically flagged and generate tickets via Pure1 alert rules.
- Automated ticket routing reduces manual intervention time.
- Capacity is actively managed against forecasts using Everpure Fusion and Pure1 capacity planners.

Everpure enablers:

- Pure1 insights engine
- Everpure Fusion and Pure1 capacity planners
- Pure1 alert rules and alert correlation tooling
- Pure1 ticket routing model
- Everpure Fusion SLA severity index

Maternity path: Level 3 → Level 4

At ML-4, repeatable tasks are automated via playbooks and monitoring tools surface trends and anomalies for intervention.

Category	Requirement
Technology enhancements	Enable predictive analytics for anomaly detection and automated resource tuning.
Process maturity	Shift operations toward exception-based handling; implement continuous improvement of operations.
Governance and measurement	Operational KPIs are shared across engineering, SRE, and infrastructure teams, with executive oversight of trends.
Benefits and metrics	Performance issues are prevented before user impact; operations drive efficiency gains.
Everpure features	Pure1 predictive engine, Everpure Fusion trend modeling, and exception-handling triggers

Technology enhancements

Requirement: Enable predictive analytics for anomaly detection and automated resource tuning.

Recommended actions:

- Deploy predictive modeling to forecast capacity and performance needs before thresholds are reached.
- Integrate the Pure1 predictive engine with an internal workload profiling data lake to consume historical application I/O signatures, enabling adaptive resource forecasting and right-sizing recommendations.
- Implement Level 1 closed-loop automation for resource scaling (for example, volume resizing) via the Everpure Fusion policy engine, where automated actions are triggered by Pure1 predictive models and require a human approval gateway.
- Integrate trend modeling with inventory to identify consolidation and optimization projects proactively.
- Configure exception-handling triggers to manage and escalate events falling outside routine automation.

Who to involve:

- **Infrastructure architect:** Defines the architecture to integrate trend modeling with inventory for proactive consolidation and optimization project identification
- **Automation engineer:** Implements the automation for nondisruptive resource tuning, volume resizing, and self-service portals
- **Virtualization engineer:** Provides critical feedback to the continuous integration loop regarding the effectiveness of automated resource tuning for virtual machine workloads
- **Data scientist/ML engineer:** Develops operations and fine-tunes the predictive models used for capacity and performance forecasting, ensuring accuracy in anomaly detection and resource tuning recommendations

Success indicators:

- Predictive engine forecasts lead to preemptive capacity increases, preventing performance bottlenecks.
- Automated volume resizing reduces the need for manual resource allocation requests.
- Exception-handling triggers ensure high-risk events are escalated immediately for manual review and possible overrides.
- Budgets adjust dynamically based on performance and efficiency (at the ML-4 state).

Everpure enablers:

- Pure1 predictive engine
- Pure1 trend modeling and Everpure Fusion workload recommendations
- Exception-handling triggers
- Everpure Fusion policy engine
- Pure1 model pipeline hooks
- Portworx® Autopilot scaling policies

Process maturity

Requirement: Shift operations toward exception-based handling; implement continuous improvement operations.

Recommended actions:

- Establish self-service portals and standardized presets to automate repeatable provisioning tasks.
- Formalize the continuous improvement loop through a structured exception review board to codify patterns derived from exception-handling logs, continually expanding the scope of autonomous processes.
- Implement exception-handling rules to manage events that fall outside of routine automation.

Who to involve:

- **Automation engineer:** Builds and maintains the self-service provisioning portals and implements new automation rules derived from the continuous improvement loop
- **Storage administrator:** Provides input on the design of standardized self-service presets and performs manual review and override for events caught by exception-handling rules
- **IT service manager:** Manages the service catalog and self-service portal, ensuring the automation rules meet user requirements and align with enterprise service management strategy

Success indicators:

- The percentage of operational staff time dedicated to reactive tasks falls below a defined threshold.
- Continuous improvement cycles generate quantifiable efficiency gains quarter over quarter.
- >75% of new deployments use standardized, automated Everpure Fusion presets (either prescriptive or customizable).

Everpure enablers:

- Exception-handling triggers
- API-first automation for self-service workflows
- Pure1 capacity planner and Everpure Fusion workload planner for continuous improvement review data

Governance and measurement

Requirement: Operational KPIs are shared across engineering, SRE, and infrastructure teams, with executive oversight of trends.

Recommended actions:

- Align operational KPIs (TTR, utilization, automation rate) across functional engineering and SRE teams.
- Integrate efficiency KPIs into executive-level dashboards for trend oversight and decision-making.
- Introduce policy-based controls such as quotas and approval workflows for high-cost or high-risk tiers.

Who to involve:

- **Enterprise application architect:** Ensures that policy-based controls (quotas and approval workflows) and efficiency KPIs align with application roadmaps and business unit objectives
- **CIO/CTO:** Provides executive oversight of operational trends and champions the integration of efficiency KPIs into executive-level dashboards for strategic decision-making
- **Procurement:** Collaborates on defining and implementing policy-based controls (quotas, approval workflows) for high-cost or high-risk tiers to ensure budget compliance and vendor management
- **Head of Infrastructure:** Ensures that operational KPIs (TTR, utilization, automation rate) are standardized and aligned across engineering, SRE, and infrastructure teams

Success indicators:

- Executive reviews are driven by unified operational trend data.
- New high-cost workloads are automatically subjected to approval workflows.
- Operational KPIs (for example, automation rate) become integral to cross-functional team OKRs.

Everpure enablers:

- Pure1 model pipeline hooks for trend analysis
- Everpure Fusion and Pure1 financial correlation view for cost impact
- Everpure Fusion policy engine for policy-based controls

Benefits and metrics

Requirement: Performance issues are prevented before user impact; operations drive efficiency gains.

Recommended actions:

- Measure the realized cost savings from identified and executed consolidation/optimization projects.
- Quantify the number of performance issues preempted by predictive analytics.
- Track data center energy and resource consumption reduction against the 20% target.

Who to involve:

- **Financial analyst/CFO:** Measures the realized cost savings from consolidation projects and ensures efficiency gains (for example, energy use reduction) are quantified in financial reporting
- **Site reliability engineer:** Quantifies the number of performance issues preempted by the predictive analytics engine to prove the value of proactive monitoring

Success indicators:

- The percentage of capacity overprovisioning is reduced below a target threshold.
- A measurable reduction in operational expense is tied directly to automation efforts.
- Efficiency gains (for example, utilization, energy use) are recognized in executive financial reporting.

Everpure enablers:

- Pure1 predictive engine for quantifying preempted issues
- Pure1 volume metrics for utilization and consolidation tracking
- Pure1 model pipeline hooks for integrating trend data into financial models

Everpure features

Requirement: Pure1 predictive engine, Everpure Fusion trend modeling, and exception-handling triggers

Recommended actions:

- Leverage the Pure1 predictive engine and Everpure Fusion trend modeling to proactively forecast capacity and performance needs.
- Configure exception-handling triggers to manage and escalate events that fall outside of routine automation.
- Use the Everpure Fusion policy engine to govern resource allocation and automated tuning processes.

Who to involve:

- **Infrastructure architect:** Configures predictive engine and trend modeling to proactively forecast capacity and performance needs, linking them to proactive project identification
- **Automation engineer:** Uses the Everpure Fusion policy engine to govern resource allocation and automated tuning processes, minimizing manual intervention

Success indicators:

- Predictive engine forecasts result in preemptive capacity increases, successfully preventing performance bottlenecks.
- Exception-handling triggers are effective, immediately escalating high-risk events for manual review.
- Trend modeling proactively identifies consolidation and optimization projects before issues arise.

Everpure enablers:

- Pure1 predictive engine
- Everpure Fusion and Pure1 trend modeling
- Exception-handling triggers
- Everpure Fusion policy engine
- Pure1 model pipeline hooks

Maternity path: Level 4 → Level 5

ML-5 is the pinnacle state, where AI-based recommendations improve performance and reduce overprovisioning, while storage tiering and provisioning are policy-guided.

Category	Requirement
Technology enhancements	Adopt AI-assisted auto-remediation and real-time workload rebalancing infrastructure.
Process maturity	Run autonomous remediation for common issues; apply intent-based operations across platforms.
Governance and measurement	Efficiency and availability are embedded into OKRs across business units; governance includes real-time operations.
Benefits and metrics	Downtime risks are near zero; operational outcomes are forecasted and preemptively optimized.
Everpure features	Everpure Fusion AI remediation modules, Pure1 intent-based recommendations, and real-time operations scoring

Technology enhancements

Requirement: Adopt AI-assisted auto-remediation and real-time workload rebalancing infrastructures.

Recommended actions:

- Implement autonomous, intent-based policies for common operational issues and resource provisioning.
- Deploy real-time workload rebalancing infrastructure to preemptively optimize performance and cost.
- Achieve intent-based orchestration by deploying Everpure Fusion AI remediation modules directly into the Everpure Intelligent Control Plane, allowing policies to declare desired state (the “intent”) and letting the system autonomously execute and self-heal deviations.

Who to involve:

- **Enterprise application architect:** Defines and implements autonomous, intent-based policies that govern operational issues and resource provisioning
- **Site reliability engineer:** Audits autonomous systems, validates AI/ML model outputs, and measures the success rate and frequency of autonomous remediation actions
- **Head of Infrastructure/CIO:** Embeds efficiency and availability metrics directly into business unit OKRs and governs cost-to-serve and unit economics.
- **Data scientist/ML engineer:** Integrates AI/ML models into the control plane for self-healing and auto-remediation, ensuring model integrity and performance

Success indicators:

- Common operational issues are resolved autonomously without human intervention.
- Workload utilization and balancing achieve near-perfect optimization.
- The platform delivers real-time operations scoring and intent-based recommendations.

Everpure enablers:

- Everpure Fusion AI remediation modules
- Pure1 intent-based recommendations
- Everpure Fusion policy engine for intent-based operations

Process maturity

Requirement: Run autonomous remediation for common issues; apply intent-based operations across platforms.

Recommended actions:

- Implement a chaos engineering discipline against the Everpure Unified Data Plane, running controlled, periodic disaster injection tests to validate the autonomous system's resilience and self-healing logic in real time.
- Define and implement intent-based operations using standardized policies (for example, service catalog provisioning).
- Implement a full chargeback model to establish storage as a financial service, tracking usage, cost-to-serve, and unit economics at the business unit level.

Who to involve:

- **Site reliability engineer:** Leads the auditing of autonomous systems and the execution of periodic, simulated failure/disaster injection tests to validate system resilience
- **Storage architect/storage engineer/administrator:** Defines the intent-based policies and service catalog structure that the autonomous remediation modules will utilize for provisioning and management
- **Financial analyst/CFO:** Designs, implements, and governs the full chargeback model, defining the calculation for cost-to-serve and unit economics to establish storage as a financial service

Success indicators:

- Autonomous remediation eliminates a significant percentage of manual Level 1 and Level 2 tickets.
- New workload provisioning is initiated and completed via a fully automated, intent-based service catalog.
- Storage is run as a financial service, with chargeback/showback models tracking usage (ML-5 state).

Everpure enablers:

- Everpure Fusion AI remediation modules
- Pure1 intent-based recommendations
- Evergreen//One adjustable service tiers for financial service alignment

Governance and measurement

Requirement: Efficiency and availability are embedded into OKRs across business units; governance includes real-time operations.

Recommended actions:

- Embed storage efficiency and availability metrics directly into business unit OKRs.
- Implement real-time governance that continuously validates compliance with intent-based policies.
- Include cost-to-serve, unit economics, and forecast variance in CIO metrics.
- Integrate real-time operations scoring directly into business unit dashboards for continuous compliance validation and rapid decision-making.

Who to involve:

- **CIO/CTO:** Embeds storage efficiency and availability OKRs into business unit performance and monitors cost-to-serve and forecast variance
- **Application owner/LoB lead:** Holds accountability for the efficiency and availability of OKRs embedded in business unit performance metrics, driving consumption alignment with business needs
- **Financial analyst:** Compiles and validates financial metrics (cost-to-serve, unit economics, forecast variance) and presents them to executive leadership for strategic decision-making

Success indicators:

- Governance reporting is based on real-time operational data.
- Efficiency becomes a core component of business unit performance measurement.
- Cost-to-serve is consistently tracked and used for financial planning.

Everpure enablers:

- Real-time operations scoring
- Evergreen//One dashboards for cost-to-serve and unit economics
- Pure1 metadata scanning for real-time compliance validation

Benefits and metrics

Requirement: Downtime risks are near zero; operational outcomes are forecasted and preemptively optimized.

Recommended actions:

- Measure efficiency by tracking the elimination rate of operational backlog (for example, ticket backlog elimination rate).
- Track the mean time to autonomous resolution for in-scope events, reporting the metric to the CIO as the definitive measure of AI/ML operational success and the elimination of Tier 1/2 ticket backlogs.
- Quantify cost-to-serve and unit economics across different business units.
- Report on the success rate and frequency of autonomous remediation actions.

Who to involve:

- **Financial analyst/CFO:** Quantifies cost-to-serve and unit economics across business units and minimizes forecasting variance through accurate financial modeling
- **Site reliability engineer:** Tracks and reports the success rate and frequency of autonomous remediation actions and the backlog elimination rate to demonstrate pinnacle operational maturity

Success indicators:

- Forecasting variance is minimized, leading to accurate budget modeling.
- Operational outcomes are preemptively optimized by AI-assisted systems.
- Downtime risks approach zero due to autonomous self-healing capabilities.

Everpure enablers:

- Everpure Fusion AI remediation modules for tracking autonomous action success
- Real-time operations scoring for continuous performance evaluation
- Pure1 intent-based recommendations to guide proactive optimization

Everpure features

Requirement: Everpure Fusion AI remediation modules, Pure1 intent-based recommendations, and real-time operations scoring

Recommended actions:

- Deploy Everpure Fusion AI remediation modules to autonomously resolve common operational issues.
- Implement Pure1 intent-based recommendations to guide and validate automated provisioning and optimization actions.
- Utilize real-time operations scoring to provide continuous, high-fidelity metrics for executive governance and audit.
- Quantify cost-to-serve and unit economics across different business units.
- Report on the success rate and frequency of autonomous remediation actions.

Who to involve:

- **Site reliability engineer:** Audits and validates the performance of Everpure Fusion AI remediation modules and ensures real-time operations scoring is integrated for governance
- **Financial analyst/CFO:** Defines the financial metrics (cost-to-serve, unit economics) and ensures they are accurately tracked and reported using Evergreen//One dashboards for business unit alignment

Success indicators:

- Common operational issues are resolved without human intervention, driven by Everpure Fusion AI remediation modules.
- Real-time operations scoring is integrated into governance and performance dashboards.
- Intent-based recommendations guide continuous optimization, minimizing overprovisioning.
- Forecasting variance is minimized, leading to accurate budget modeling.
- Operational outcomes are preemptively optimized by AI-assisted systems.
- Downtime risks approach zero due to autonomous self-healing capabilities.

Everpure enablers:

- Everpure Fusion AI remediation modules
- Pure1 intent-based recommendations
- Real-time operations scoring
- Everpure Fusion policy engine
- Evergreen//One adjustable service tiers

Maturity path: Level 5

Finally, operations run with a supervised, autonomous execution across storage domains, while CIOs monitor SLO adherence, auto-remediation frequency, and backlog elimination rates.

The autonomous engine of efficiency

The journey to ML-5 marks a critical shift from manually managing infrastructure to running a supervised, autonomous EDC. This pinnacle state is where the strategic vision of increased efficiency is fully realized, transforming storage from a reactive cost center into a continuously optimized financial service. By embedding the Intelligent Control Plane properties of Everpure Fusion AI remediation capabilities and Pure1 predictive intelligence directly into the Unified Data Plane, operational expenditure is slashed by the target 74%. The platform achieves near-perfect resource optimization, evidenced by an 85% utilization rate across the storage fleet. The focus of operational teams pivots entirely: human effort shifts from execution and firefighting to auditing autonomous systems and driving business unit OKRs through real-time efficiency metrics. Ultimately, this self-healing, policy-governed environment ensures that efficiency is no longer an aspiration but a sustained, measurable competitive advantage.

Chapter 2: Business continuity

Establishing an actionable framework for business continuity is paramount to minimizing an organization's exposure to operational and financial risks. The business continuity pillar in this blueprint focuses on moving enterprises away from fragmented, reactive emergency responses toward a proactive, unified posture.

Progressing through business continuity maturity levels transforms disaster recovery from a siloed, manual process into an automated, self-healing operational asset, ultimately securing the path toward an enterprise architecture that guarantees 99.9999% reliability and defends long-term business resilience.

This chapter outlines the path from manual recovery to a self-healing, autonomous resilience posture within the risk exposure reduction outcome.

Pillar overview: Strategic impact and metrics

- **Impact:** Keep critical services running through infrastructure failures or external events.
- **Target metric:** 99.9999% reliability
- **Library focus:** Disaster recovery, failover automation, resilience governance, and SLA guarantees

Maturity level introduction

Most enterprises don't set out to build fragmented continuity strategies. They build application by application, team by team, and end up with a patchwork of protection policies that nobody fully owns or has fully tested. That's Level 1. It's more common than most organizations admit.

The following five maturity levels describe a progression from that starting point to infrastructure that detects, responds to, and recovers from disruption without waiting for human intervention. Each level is a meaningful operational state, not just a checkpoint. Organizations at Level 3 look fundamentally different from organizations at Level 2, and the gap between Level 4 and Level 5 is the difference between fast recovery and little to no business impact.

Most organizations land somewhere between Level 2 and Level 3. Getting past Level 3 is where progress tends to slow not because the technology is hard, but because the next steps require process and governance changes, not just new tooling. This blueprint is designed to unblock exactly that.

Maturity levels	
Level	Maturity description
ML-1	• There is no documented business continuity strategy, and workload protection is inconsistent. • Recovery timelines are undefined or unmanaged.
ML-2	• Backup schedules and recovery procedures are defined for critical systems and tested annually. • Recovery time objective (RTO) and recovery point objective (RPO) targets exist but vary widely by team.
ML-3	• Continuity plans are operationalized and tracked. • Automated failover exists for mission-critical applications, and restoration testing occurs quarterly.
ML-4	• End-to-end orchestration with failover processes is integrated with real-time monitoring. • Simulation exercises validate business impact and response.
ML-5	• Resilience is self-healing, as systems detect disruption and reroute services autonomously. • CIO reporting includes verified failover times, recovery KPIs, and compliance with disaster recovery policies across hybrid environments.

Maturity path summary

The four maturity transitions in the following table represent the full arc from manual, undocumented recovery to self-healing, board-reported resilience. Each row is a condensed view of what needs to change across technology, process, and governance, along with the Everpure capabilities that enable movement from one level to the next. Think of it as the map before you read the directions.

There are three things worth noting before you use it. First, the sequencing is deliberate. The technology changes at Level 3 assume the process foundation from Level 2 is in place. Skip the ownership assignment, documented playbooks, and automated failover surfaces for every gap you don't close. Second, you don't have to start at Level 1. If your organization is already operating at Level 2 or Level 3 in specific areas, start there. The blueprint can be entered at any point. Third, not every column moves at the same pace. Some organizations have strong governance but weak technology. Others have the tooling but none of the process accountability around it. The table helps you spot the gaps, not just where to go next.

Paths				
Maturity transition	Technology enhancements	Process maturity	Governance	Enabling features
Level 1 → 2	Deploy local backup and snapshot infrastructure with basic recovery capabilities.	Document recovery procedures for critical systems; perform manual failover drills periodically.	Assign recovery ownership by application; IT team is responsible for continuity documentation and testing.	SafeMode™ Snapshots, local snapshot scheduling in Everpure Fusion, and manual recovery orchestration
Level 2 → 3	Implement automated failover systems with cross-site replication and remote site readiness.	Standardize disaster recovery plans with role-specific playbooks and maintenance schedules; automate test validations.	Create a dedicated disaster recovery coordination team; incorporate recovery metrics into operational reporting.	ActiveDR™ for replication, Everpure Fusion disaster recovery readiness workflows, and Pure1 failover validation reports
Level 3 → 4	Enable application-aware recovery with orchestrated workflows across hybrid environments.	Continuity plans are embedded in the continuous innovation/continuous delivery (CI/CD) life cycle; drills are integrated with change and release processes.	Resilience targets are governed by a cross-functional steering committee; risk is reviewed across business units.	Everpure Fusion application protection policies, Everpure Resilience service orchestration, and API-based recovery sequencing
Level 4 → 5	Adopt predictive recovery infrastructure using AI models, real-time impact scoring, and auto-remediation triggers.	Run live-swap readiness with no-downtime validation; recovery plans simulate full-stack dependencies.	Continuity is embedded into the enterprise risk model; the board and regulators review resilience KPIs.	Predictive recovery modeling via Pure1, Everpure Fusion event-based orchestration, and ActiveCluster™ high availability automation

Maturity path: Level 1 → Level 2

The move to documented recovery

In ML-1, the business continuity strategy is largely undocumented, and recovery timelines are inconsistent and untested. Resilience is dependent on manual effort.

This transition focuses on establishing predictable schedules and procedures for critical systems.

Category	Requirement
Technology enhancements	Deploy local backup and snapshot infrastructure with basic recovery capabilities.
Process maturity	Document recovery procedures for critical systems, and perform manual failover drills periodically.
Governance and measurement	Assign recovery ownership by application; IT team is responsible for continuity documentation and testing.
Benefits and metrics	Backup schedules and RTO/RPO targets are established for critical systems.
Everpure features	SafeMode protected snapshots, local snapshot scheduling in Everpure Fusion, and manual recovery orchestration

Technology enhancements

Requirement: Deploy local backup and snapshot infrastructure with basic recovery capabilities.

Recommended actions:

- Audit current storage arrays to confirm snapshot capabilities are active and not relying solely on legacy streaming backups.
- Configure cross-site replication baselines for all storage volumes hosting core operational data.
- Set up a centralized dashboard environment to gain clear, multi-array visibility into data protection and replication health status.

Who to involve:

- **Storage administrators/storage engineers:** Configures local snapshot schedules and maps out initial target replication paths
- **Infrastructure architects:** Ensures target recovery sites or cloud regions have matching capacity allocations to receive incoming replication data

Success indicators:

- **Decommissioned manual backups:** Complete retirement of manual, administrator-triggered backup routines across all core systems
- **Automated snapshot infrastructure:** Deployment of automated local snapshot policies active across 100% of core workloads
- **Immutable data protection:** Enforcement of tamper-proof data locking via configurations like SafeMode Snapshots to block accidental or malicious administrative deletion

Everpure enablers:

- SafeMode protected snapshots
- ActiveDR replication

Process maturity

Requirement: Document recovery procedures for critical systems and perform manual failover drills periodically.

Recommended actions:

- Facilitate workshops with business units to categorize applications by criticality and clearly define RTO and RPO targets.
- Map out and document sequential, step-by-step disaster recovery workflows into an official playbook, eliminating dependency on institutional knowledge. **Note:** This playbook should also be available offline to ensure access during network outages or disasters.
- Establish a locked-in monthly operational calendar window for manual backup verification and restore tests.

Who to involve:

- **Application owners and product managers:** Quantify maximum tolerable downtime data points and sign off on realistic RTO/RPO targets
- **Business continuity/disaster recovery managers:** Draft the new standardized playbooks and govern their execution
- **Backup and recovery specialists:** Run and log the mandatory monthly recovery validation tests

Success indicators:

- **Standardized recovery playbooks:** Completion of written, step-by-step restoration documentation for every critical application, eliminating dependency on institutional knowledge
- **Independent execution capability:** Verification that any certified infrastructure administrator can execute a successful restore using playbooks alone
- **Locked-in testing rhythm:** Transition of recovery testing into a mandatory monthly calendar event with documented verification logs

Everpure enablers:

- Pure1 and Everpure Fusion policy templates

Governance and measurement

Requirement: Assign recovery ownership by application; IT team is responsible for continuity documentation and testing.

Recommended actions:

- Formally designate an internal disaster recovery lead role within the core infrastructure operations team to serve as a single point of accountability.
- Build automated data extraction templates to pull recovery readiness statistics monthly.
- Launch a quarterly reporting cadence to proactively present data protection compliance posture and test recovery results directly to application stakeholders.

Who to involve:

- **Infrastructure director/IT operations manager:** Assigns internal ownership roles and enforces operational accountability
- **Risk and compliance officers:** Ensure that the defined metrics align directly with industry compliance standards or corporate mandates
- **Application team leads:** Receive, review, and challenge quarterly health metrics regarding their respective platform and dependencies

Success indicators:

- **Established disaster recovery accountability:** Formal assignment of a dedicated disaster recovery lead role within the infrastructure operations team
- **Proactive compliance reporting:** Implementation of a predictable quarterly reporting schedule to push data protection health metrics directly to application stakeholders

Everpure enablers:

- Pure1 role-based access controls
- Pure1 dashboards
- Pure1 data protection assessment

Benefits and metrics

Requirement: Backup schedules and RTO/RPO targets are established for critical systems.

Recommended actions:

- Establish a baseline for the time to discover and time to remediate snapshot and backup execution errors.
- Track the percentage of critical workloads successfully covered under formalized, documented playbooks versus those relying on legacy ad hoc recovery strategies.
- Measure and report on your backup/recovery testing success rate (target 100% of scheduled monthly tests executed and validated).

Who to involve:

- **CIO/CTO:** Uses foundational baseline metrics as an objective validation of risk mitigation progression
- **IT operations analytics:** Tracks, compiles, and verifies the reduction in recovery delays and playbook validation gaps

Success indicators:

- **Tier target alignment:** Documented business unit alignment and sign-off on explicit, binding RTO and RPO metrics
- **On-demand posture verification:** Instant access to dashboard telemetry to mathematically prove active data loss windows and exposure risk limits at any given moment

Everpure enablers:

- Pure1 snapshot visibility dashboards
- Pure1 telemetry

Everpure features

Requirement: Asynchronous replication policy enforcement, SafeMode Snapshot protection across all arrays, and centralized Pure1 visibility to validate RPO conformance and close replication gaps before an incident occurs

Recommended actions:

- Provision and apply asynchronous replication policies across storage architectures to secure required RPO objectives.
- Activate SafeMode protection across all arrays to ensure snapshots are completely immutable and indelible from accidental or malicious deletion.
- Utilize centralized Pure1 snapshot visibility to track overall retention policy conformance and catch data replication gaps before an incident occurs.

Who to involve:

- **Dedicated systems/storage engineers:** Activate configurations and maintain policy compliance on live storage systems
- **Everpure account team/technical account manager (TAM):** Review SafeMode protected snapshot configuration best practices, verify snapshot lock periods, and ensure architecture layouts match best practices

Success indicators:

- **Consolidated management control:** Reduction in operational friction by shifting primary backup verification from fragmented third-party software tools directly into the native platform interface
- **Frontline recovery capability:** Verification that Tier 1, on-call operations personnel can successfully trigger and verify local snapshot restores without needing escalation to specialized storage engineers
- **Zero overhead baseline:** Complete elimination of infrastructure performance degradation or application latency spikes during snapshot creation and local data restoration windows

Everpure enablers:

- Everpure SafeMode protected snapshots
- Pure1 data protection assessment

Maturity path: Level 2 → Level 3

The move to automated failover

Transitioning to ML-3 operationalizes continuity plans by introducing automated failover for mission-critical applications and standardized testing cycles.

Category	Requirement
Technology enhancements	Implement automated failover systems with cross-site replication and remote site readiness.
Process maturity	Standardize disaster recovery plans with role-specific playbooks and maintenance schedules; automate test validations.
Governance and measurement	Create a dedicated disaster recovery coordination team; incorporate recovery metrics into operational reporting.
Benefits and metrics	Continuity plans are operationalized; quarterly restoration testing is achieved for key applications.
Everpure features	ActiveDR for replication, Everpure Fusion disaster recovery readiness workflows, and Pure1 failover validation reports

Technology enhancements

Requirement: Implement automated failover systems with cross-site replication and remote site readiness.

Recommended actions:

- Group organizational workloads into explicit protection tiers (for example, Tier 0 Critical, Highly Privileged Systems; Tier 1 Mission-Critical; Tier 2 Business-Critical) within the storage fleet rather than treating all storage arrays uniformly.
- Integrate automated recovery scoring mechanisms that continually parse active data replication lag and calculate real-time RTO/RPO readiness state.

Who to involve:

- **Storage architects/infrastructure engineers:** Map data volumes to policy tiers and establish automated, multisite replication orchestration
- **Enterprise platform administrators:** Configure tracking APIs and ensure infrastructure telemetry feeds directly into your monitoring systems

Success indicators:

- Systematic grouping of distinct system workloads into defined, tier-based application protection matrices
- Configuration of matching, remote site infrastructure targets verified as ready-to-accept, automated data payloads

Everpure enablers:

- Everpure ActiveCluster
- Everpure ActiveDR

Process maturity

Requirement: Standardize disaster recovery plans with role-specific playbooks and maintenance schedules; automate test validations.

Recommended actions:

- Standardize and institutionalize a quarterly calendar to perform nondisruptive disaster recovery tests across critical business environments.
- Document post-test processes and outcomes, recovery timeline benchmarks, and discovered gaps.
- Implement automated alerting to enforce targeted recovery SLAs across all active production environments.

Who to involve:

- **Disaster recovery coordinators/specialists:** Orchestrate, execute, and log the quarterly replication and recovery validation exercises
- **DevOps/SRE teams:** Build automatic alerting that flags when a service falls below SLA baselines
- **Application owners:** Participate in test verification and validate post-recovery data consistency

Success indicators:

- Transition of disaster recovery validation routines from annual events to a locked-in, mandatory quarterly drill schedule
- Automation of verification checks for replicated data volumes, removing manual operator testing steps
- Standardized capture and archival of post-drill logging artifacts to document timeline bottlenecks and system drift

Everpure enablers:

- Pure1 and Everpure Fusion disaster recovery workflows

Governance and measurement

Requirement: Create a dedicated disaster recovery coordination team; incorporate recovery metrics into operational reporting.

Recommended actions:

- Establish a dedicated, cross-functional disaster recovery governance committee to meet regularly, review readiness metrics, and govern any remediation backlog.
- Formally assign individual risk owners within each business unit or core platform layer who are accountable for accepting, deferring, or funding data resilience gaps.

Who to involve:

- **CISO/Chief Risk Officer (CRO):** Aligns business unit risk policies with overall organizational continuity strategies and compliance goals
- **Business unit directors/application VPs:** Act as the designated system risk owners and sign off on identified tier-based protection metrics
- **IT governance managers:** Operationalize the reporting cadence and maintain risk accountability

Success indicators:

- Formal chartering of a dedicated, cross-functional disaster recovery coordination team, with active coordination meetings
- Direct embedding of recovery timeline compliance metrics into automated IT infrastructure operational reporting lines
- Formally designated and named risk owners assigned to oversee accountability within individual business units

Everpure enablers:

- Pure1 asset topology view
- Pure1 data protection assessment

Benefits and metrics

Requirement: Continuity plans are operationalized; quarterly restoration testing is achieved for key applications.

Recommended actions:

- Calculate and trend the automated failover success rate across all operational tiers.
- Measure the total reduction in recovery gaps by tracking test SLA compliance rates over consecutive quarters.

Who to involve:

- **CIO/Director of IT Operations:** Reviews automated failover performance data as evidence of reduced financial and operational risk exposure
- **IT business analysts/capacity planners:** Review log data, compile trends, and publish readiness reviews

Success indicators:

- 100% conformance showing quarterly restoration test execution metrics meet or exceed defined business SLA milestones
- Stabilized and mathematically predictable baseline metrics across all active data replication pipelines

Everpure enablers:

- Pure1 continuous replication

Everpure features

Requirement: Policy-based, multisite replication; centralized disaster recovery monitoring; and a validated secondary recovery environment for safe failover testing and operational readiness

Recommended actions:

- Leverage the Everpure Fusion engine to build out and enforce tier-based disaster recovery layout rules across your multiregion fleet.
- Monitor Pure1 dashboards to gain a single-pane-of-glass view of continuous data replication and instance compliance metrics.
- Configure cross-site failover replication using ActiveDR and environment sandbox validation inside Everpure Resilience service to safely test full-stack recovery without impacting production.

Who to involve:

- **Storage administrators:** Provision policy templates, review telemetry, and ensure configuration rules map consistently across data environments
- **Everpure field solution architects/account team:** Run architecture validation through Everpure Value Reviews, verify any automated failover scripts, and ensure best practices for nondisruptive testing

Success indicators:

- Retirement of custom, human-managed secondary array replication logic in favor of native platform automation
- Single-interface visibility over multisite recovery point consistency states, eliminating siloed infrastructure checks

Everpure enablers:

- Everpure Fusion policies
- Pure1 continuous replication view
- Everpure Resilience service

Maturity path: Level 3 → Level 4

The move to orchestrated hybrid resilience

This transition represents a critical pivot: moving from isolated, storage-level automated failovers to orchestrated hybrid resilience, where data continuity becomes completely application-aware and directly embedded into your customers' modern software delivery pipelines.

Category	Requirement
Technology enhancements	Enable application-aware recovery with orchestrated workflows across hybrid environments.
Process maturity	Embed continuity plans in the CI/CD life cycle; drills are integrated with change and release processes.
Governance and measurement	Resilience targets are governed by a cross-functional steering committee; risk is reviewed across all business units.
Benefits and metrics	Simulation exercises validate business impact; recovery is fully integrated with real-time operations.
Everpure features	Everpure Fusion application protection policies, Everpure Resilience service orchestration, and API-based recovery sequencing

Technology enhancements

Requirement: Enable application-aware recovery with orchestrated workflows across hybrid environments.

Recommended actions:

- Elevate disaster recovery groups from volume/array-level definitions to full-stack, application-aware consistency (for example, group web, middleware applications, and database volumes into a single logical unit).
- Establish automated, multiregion and hybrid-cloud recovery orchestrations that span on-premises arrays and public cloud instances seamlessly.
- Integrate application mapping tools to ensure system components spin up in the same sequence required to prevent data corruption.

Who to involve:

- **Enterprise cloud architects:** Bridge on-premises architecture needs with public cloud recovery landing zones
- **Database administrators:** Architect database failover sequencing and integrity checks, ensuring that multi-tier applications spin up in the correct order and resume operations across sites without data corruption or transaction loss
- **Infrastructure automation engineers:** Write and maintain infrastructure-as-code orchestration playbooks required to spin up, network, and validate full application stacks automatically across hybrid or multicloud recovery zones.

Success indicators:

- Transition from raw, volume-level storage group definitions to complex, application-aware consistency boundaries.
- Establish multiregion and hybrid-cloud recovery orchestrations that seamlessly map data between on-premises and cloud targets.

Everpure enablers:

- Everpure Resilience service

Process maturity

Requirement: Embed continuity plans in the CI/CD life cycle; drills are integrated with change and release processes.

Recommended actions:

- Shift data resilience into software engineering workflows by managing protection parameters using infrastructure-as-code models.
- Embed automated, isolated sandbox disaster recovery validation testing directly inside the CI/CD pipelines.
- Transition from scheduled, choreographed testing windows to regular, unannounced failover simulation drills executed inside isolated network environments during regular business hours.

Who to involve:

- **DevOps/platform engineers:** Write and embed disaster recovery policy confirmations directly into pipelines (for example, GitHub, Terraform)
- **Release managers:** Govern code deployment quality and mandate updated failover definitions
- **Site reliability engineers:** Design and execute unannounced, nondisruptive sandbox simulations

Success indicators:

- Permanent inclusion of data resilience and disaster recovery configurations within automated deployment codebases
- Execution of unannounced, nondisruptive, partial-failover simulation drills inside isolated network segments during standard business hours

Everpure enablers:

- Everpure Resilience service

Governance and measurement

Requirement: Resilience targets are governed by a cross-functional steering committee; risk is reviewed across all business units.

Recommended actions:

- Transition from basic operations review team to a formalized, cross-functional resilience steering committee, which includes decision-makers from engineering, business units, legal, security, and finance.
- Map technical recovery metrics directly to business outcomes, ensuring that infrastructure capabilities prevent financial penalties or breaches of customer-facing SLAs.
- Standardize application programmatic remediation and force code fixes for any application that fails resilience validation.

Who to involve:

- **VP of Engineering/VP of Infrastructure:** Sponsors and runs the cross-functional resilience steering committee
- **Corporate legal/compliance teams:** Tie storage resilience metrics directly to regulatory obligations and client liabilities
- **Product management leadership:** Prioritizes engineering cycles to address identified business continuity gaps

Success indicators:

- Formal operational tracking run by an active, cross-functional resilience steering committee
- Direct mathematical alignment of technical data replication metrics with contractual, customer-facing SLAs

Everpure enablers:

- Pure1 REST APIs

Benefits and metrics

Requirement: Simulation exercises validate business impact; recovery is fully integrated with real-time operations.

Recommended actions:

- Measure disaster recovery-as-code coverage rate (target 100% of newly deployed applications to feature native recovery definitions).
- Track simulation velocity, calculating the time from simulation trigger to end-to-end data verification.
- Quantify and report on the difference between calculated RTO/RPO and simulated disruption test metrics.

Who to involve:

- **CTO/CIO:** Leverages high-level metrics to prove a minimized risk to executive boards
- **FinOps/IT financial analysts:** Model the money saved by reducing deployment recovery validation

Success indicators:

- 100% of infrastructure deployments automatically inherit mandated corporate data protection playbooks
- Documented reduction in verification velocities, calculating the speed from simulation trigger to multi-tier consistency confirmation

Everpure enablers:

- Pure1 and Everpure Fusion disaster recovery workflows

Everpure features

Requirement: Application-aware protection policies, orchestrated recovery across on-premises and hybrid environments, and API-driven integration with operational alerting and service workflows

Recommended actions:

- Implement Everpure Fusion protection policies to package multi-volume application components into a single, protected logical bundle that adheres to implemented policies.
- Deploy Everpure Resilience service to orchestrate on-premises or hybrid-cloud recovery, allowing for multisite, application-aware executions.
- Connect to exposed Pure1 REST APIs to stream live data protection status and compliance alerts directly to tooling (for example, ServiceNow, Jira) for real-time alerting and service ticket routing.

Who to involve:

- **Senior systems integrators/DevOps engineers:** Map Pure1 API webhooks into internal alerting and ticketing software
- **Everpure customer success engineer/TAM:** Reviews Everpure Resilience service orchestration and evaluates against operational goals

Success indicators:

- Complete transition from siloed, infrastructure-level monitoring to comprehensive, workload-centric recovery frameworks
- Elimination of manual script maintenance for application sequencing through unified platform configurations

Everpure enablers:

- Pure1 application workflows
- Everpure Resilience service

Maturity path: Level 4 → Level 5

The move to autonomous, self-healing resilience

ML-5 is the pinnacle state, where resilience is self-healing, systems autonomously reroute services during disruption, and recovery metrics are fully transparent to leadership.

At Level 5, the infrastructure continuously monitors, tests, and heals itself without waiting for human intervention.

Category	Requirement
Technology enhancements	Adopt predictive recovery infrastructure using AI models, real-time impact scoring, and auto-remediation triggers.
Process maturity	Run live-swap readiness with no-downtime validation; recovery plans simulate full-stack dependencies.
Governance and measurement	Continuity is embedded into the enterprise risk model; CIO reporting includes verified failover times and hybrid disaster recovery compliance.
Benefits and metrics	Operations are self-healing; there is board-level assurance of six-nines (99.9999%) reliability across hybrid environments.
Everpure features	Predictive recovery modeling via Pure1, Everpure Fusion event-based orchestration, and ActiveCluster high availability automation

Technology enhancements

Requirement: Adopt predictive recovery infrastructure using AI models, real-time impact scoring, and auto-remediation triggers.

Recommended actions:

- Deploy AI-driven predictive systems that flag anomalies (for example, subtle latency spikes or unusual read/write patterns) and isolate potential infrastructure failures before they cause an outage.
- Configure intelligent, event-based routing layers that dynamically shift user workloads to secondary data centers when the health of a primary site degrades below acceptable thresholds.
- Standardize cross-platform automation hooks that enable real-time recovery workflows.

Who to involve:

- **AIOps/platform analytics engineers:** Build and monitor the data pipelines that ingest and process live data center telemetry.
- **Cloud and network architects:** Implement automated, dynamic traffic-routing systems that reroute active user traffic seamlessly.

Success indicators:

- Deployment of intelligent, event-driven routing layers that automatically shift active user workloads to secondary zones before hardware or site degradation leads to a hard crash
- Integration of predictive monitoring architectures that dynamically isolate failing components based on micro-anomaly telemetry signatures

Everpure enablers:

- Pure1 REST APIs
- Everpure Resilience service

Process maturity

Requirement: Run live-swap readiness with no-downtime validation; recovery plans simulate full-stack dependencies.

Recommended actions:

- Adopt advanced chaos engineering frameworks, injecting synthetic failures, automated site disconnects, and performance bottlenecks into the live production environment to prove self-healing resilience under stress.
- Evolve testing into a continuous live-swap standard, where applications move back and forth between data centers with zero downtime and zero impact to the end user.
- Implement real-time, automated application mapping to ensure recovery playbooks update whenever code changes.

Who to involve:

- **Director of SRE:** Champions and governs the chaos engineering policies and live-swap testing frameworks
- **Principal software engineers/DevSecOps leads:** Architect application codebases to be stateless or cloud-native, allowing them to survive automated, multisite routing changes

Success indicators:

- Transition to continuous live-swap operational standards, allowing production applications to float fluidly between sites with zero end user impact or downtime.
- Institutionalize automated, continuous validation checks that eliminate scheduled disaster recovery test weekends entirely.

Everpure enablers:

- Everpure Fusion AI remediation modules
- Pure1 intent-based recommendations
- Everpure Fusion policy engine for intent-based operations

Governance and measurement

Requirement: Continuity is embedded into the enterprise risk model; CIO reporting includes verified failover times and hybrid disaster recovery compliance.

Recommended actions:

- Embed data resilience metrics directly into the enterprise risk management (ERM) model to raise continuity data to the same governance tier as corporate finance, cybersecurity liability, and legal compliance.
- Deliver continuous, automated, audit-ready evidence feeds to third-party regulators, insurance underwriters, and enterprise clients, eliminating manual compliance reporting overhead.
- Establish real-time risk-tolerance thresholds, where the system automatically self-corrects if a business unit breaches its assigned risk threshold.

Who to involve:

- **CRO/board risk committee:** Governs data resilience and a pillar of corporate health
- **Regulatory compliance director:** Takes real-time, automated resilience stream into regulatory alignment certificates (for example, SEC, HIPAA, GDPR compliance)

Success indicators:

- Structural embedding of data resilience matrices and availability metrics directly into the master ERM model
- Implementation of continuous, automated data streams, providing real-time evidence of audit readiness to external regulators, insurers, and clients

Everpure enablers:

- Pure1 REST APIs
- Pure1 assessment dashboards

Benefits and metrics

Requirement: Operations are self-healing; there is board-level assurance of six-nines (99.9999%) reliability across hybrid environments.

Recommended actions:

- Track and verify the six-nines reliability metric (99.9999% availability) across the entire core infrastructure estate, allowing for no more than 31.5 seconds of total downtime per year.
- Measure the zero-touch recovery rate—the percentage of operational anomalies and site disruptions successfully resolved by automation without human intervention (target is 100%).
- Quantify the reduction in corporate compliance and audit preparation achieved by replacing manual validation reviews with automated compliance feeds.

Who to involve:

- **CIO, CEO, and executive board:** Utilize automated resilience scorecards to increase market share in highly regulated industries and assure shareholders of total business continuity

Success indicators:

- Mathematical verification of six-nines availability (99.9999% uptime) is achieved across all core enterprise application profiles.
- A 100% rating is achieved on the zero-touch recovery rate metric, indicating all infrastructure anomalies are autonomously resolved without human triage.

Everpure enablers:

- Everpure Fusion AI remediation modules
- Pure1 intent-based recommendations

Everpure features

Requirement: Predictive analytics, event-driven policy automation, and active-active multisite resilience built on a standardized, software-defined control plane

Recommended actions:

- Leverage Pure1 modeling and analytics to forecast long-term capacity, performance, and hardware health trends while running automated risk simulations across storage footprints.
- Operationalize Pure1 workflows for event-based orchestration to roll out, inherit, and lock down data protection policies across new storage arrays the moment they register on the network.
- Leverage ActiveCluster automation to run synchronous, multisite, active-active storage clustering to deliver zero-RTO/RPO capabilities across physical distances.

Who to involve:

- **Enterprise infrastructure director:** Aligns the organization's overall self-healing vision with native automation capabilities inside the Everpure hardware and software stacks
- **Everpure principal technology strategists:** Act as strategic design partners to ensure the client's automation engines ingest and trigger Everpure APIs according to architectural design blueprints

Success indicators:

- Total transition to an entirely software-defined data plane, where application continuity is decoupled from localized infrastructure boundaries
- Elimination of manual policy propagation overhead when expanding or provisioning new infrastructure components globally

Everpure enablers:

- ActiveCluster automation
- Everpure Fusion event-based policy orchestration

Maturity path: Level 5

Achieving world-class business continuity is a continuous operational evolution. By leveraging this blueprint as a self-assessment tool, organizations can chart their current position, help eliminate blind spots, and break down organizational silos that have historically stalled disaster recovery efforts. Moving from the fragmented practices of Level 1 to the autonomous, self-healing of Level 5 requires collaboration across every layer of the organization. Embracing this maturity path does more than just protect an organization from worst-case scenarios. It transforms resilience into a competitive advantage, reducing risk exposure and ensuring the organization can confidently deliver on its promise of operational reliability.

Chapter 3: Application agility and automation

Scaling with agility reflects an organization's ability to adapt infrastructure, data, and execution workflows in near-real time based on business intent, workload demand, and external change. It encompasses platform observability, workload intelligence, and the automation needed to continuously align operations with outcomes.

Without agility and signal-driven execution, platforms become bottlenecks—not enablers—of transformation. Decisions lag behind reality, capacity is misallocated, and innovation becomes gated by infrastructure limitations.

Platforms should respond to business signals in real time; telemetry informs product, go-to-market (GTM), and resource allocation decisions. Workloads are dynamically placed, tuned, and retired based on changing conditions and intent, with each step of the process aligned to established SLAs for performance, protection, and availability based on the workloads' criticality to the business.

In short, infrastructure should automatically respond to the needs of the business rather than relying on manual intervention.

Without this, execution delays, platform sprawl, underutilized infrastructure, and inability to pivot in response to market or customer changes lead to wasted spend from overprovisioning, missed automation ROI, and loss of first-mover advantage in adaptive product or GTM strategies. Meanwhile, competitors with AI-optimized, intent-aware platforms scale faster, serve customers more responsively, and pivot GTM more effectively, giving them a competitive advantage and placing your business at risk.

This chapter focuses on the application agility and automation pillar, providing action plans for moving up the maturity model in several critical areas. These actions can take your business from relying on manual processes and actions decided by humans using disjointed tools and metrics to a true, autonomous, self-driving, SLA-and-outcome-aligned EDC.

Pillar overview: Application agility and automation

- **Impact:** Full-stack application deployment is completely automated and aligned with business SLAs and outcomes, with continuous monitoring of application performance including business impact context and visualization.
- **Target metric:** On-demand application deployment measured in minutes
- **Library focus:** Data governance, SLA management, and incident resolution

Maturity level introduction

The gap between Level 1 and Level 5 in this pillar is the gap between a team that spends most of its time on tickets and a platform that resolves most issues before a ticket is ever opened. That's not an exaggeration. At Level 1, every provisioning request, performance issue, and capacity question is a manual task. At Level 5, the system handles the routine work autonomously, and the team focuses on exceptions and business outcomes.

The following five maturity levels describe that progression. Each one is a distinct operational state with its own characteristics, not just a degree of the same capability. The difference between Level 2 and Level 3 lies in the shift from reducing manual work to automating it. The difference between Level 4 and Level 5 lies in predicting problems versus eliminating them before they occur.

Most organizations are at Level 2 or are in the early stages of Level 3. The tooling exists to go further. What tends to hold teams back is the process and governance work that must accompany the technology. This blueprint is designed to unblock exactly that.

Maturity levels	
Level	Maturity description
ML-1	- Storage is managed by infrastructure teams disconnected from business cycles. - Planning is reactive and driven by incident response.
ML-2	- Cross-team visibility exists through quarterly reviews and shared planning documents. - Response to business change is slow.
ML-3	- Execution is governed by agile rituals. - Storage teams deliver capacity tied to sprint and release planning.
ML-4	- Shared metrics align storage ops to business OKRs. - Infrastructure teams pivot resources dynamically based on real-time demand.
ML-5	- Execution is fluid. - CIOs track business-aligned lead time, change responsiveness, and workload onboarding velocity.

Maturity path summary

The following table maps all four maturity transitions in a single view. Each row shows what needs to change across technology, process, and governance, along with the Everpure capabilities that enable movement from one level to the next. It's a fast way to orient yourself before stepping into the detailed guidance for each transition.

There are three things worth noting before you use it. First, the columns don't always move together. Some organizations have strong monitoring but no process accountability around what the monitoring surfaces. Others have governance structures in place but haven't connected them to the right tooling yet. The table helps you see the full picture, not just the technology gap. Second, the sequencing is intentional. The automation at Level 3 only works cleanly if the ownership and runbooks from Level 2 are already in place. Skipping steps is possible; it just tends to create technical debt that shows up later. Third, you don't have to start at Level 1. Enter at the level that reflects where your organization actually is today.

Paths				
Maturity transition	Technology enhancements	Process maturity	Governance	Enabling features
Level 1 → 2	Establish centralized task visibility, basic collaboration tools, and ticketing for ops workflows.	Introduce project tracking and functional accountability for delivery activities.	Assign execution leads per function; build recurring ops reviews into team cadences.	Everpure Fusion operational dashboards, project tracking via tags, and Pure1 delivery status telemetry
Level 2 → 3	Enable workflow automation for approvals, resource requests, and deployment tasks.	Standardize delivery runbooks; define metrics for execution predictability and team throughput.	Operations and delivery managers coordinate across departments; SLAs are embedded into execution KPIs.	Everpure Fusion automation templates, task telemetry in Pure1, and cross-team coordination controls
Level 3 → 4	Deploy adaptive planning platforms, resource simulation tools, and execution telemetry.	Align quarterly planning with delivery cycles; integrate readiness metrics into OKRs.	An execution readiness committee oversees cross-functional delivery; risk is actively tracked and managed.	Everpure Fusion project forecasting, Pure1 execution velocity analytics, and readiness scoring modules
Level 4 → 5	Adopt AI-based execution modeling, cross-pillar orchestration, and proactive risk avoidance systems.	Use predictive readiness and task dependency engines to plan and deliver adaptively.	There is board-level insight into execution constraints; scenario planning is embedded into governance cycles.	Everpure Fusion readiness engine, Pure1 adaptive risk modeling, and intelligent orchestration tools

Maturity path: Level 1 → Level 2

The move to application agility

In ML-1, applications are deployed via human activity. There is no connection between infrastructure teams and business cycles or outcomes. Planning is reactive and driven by incident response. Significant time is required to deploy new applications, and metrics are segmented and noncontextual. The infrastructure teams respond to issues at their level with little to no awareness of business impact.

Category	Requirement
Technology enhancements	Establish centralized task visibility, basic collaboration tools, and ticketing for ops workflows.
Process maturity	Introduce project tracking and functional accountability for delivery activities.
Governance and measurement	Assign execution leads per function; build recurring ops reviews into team cadences.
Benefits and metrics	Execution is often delayed or reactive; visibility is limited to siloed workflows.
Everpure features	Everpure Fusion operational dashboards, project tracking via tags, and Pure1 delivery status telemetry

Technology enhancements

Requirement: Establish centralized task visibility, basic collaboration tools, and ticketing for ops workflows.

Recommended actions:

- Deploy a single ticketing system (ServiceNow, Jira Service Management, or equivalent) as the system of record for all infrastructure work requests, replacing email and chat-based intake.
- Stand up a shared collaboration workspace (Confluence, SharePoint, or Teams) where runbooks, change records, and operational documentation live in one searchable location.
- Use Pure1 APIs to pull array health, capacity, and performance telemetry into a monitoring platform so events generate trackable work items rather than ad hoc tasks.
- Configure all storage arrays to send OpenMetrics telemetry to a centralized dashboard (Prometheus, Datadog, and so on) as the shared visual layer for real-time fleet performance and capacity metrics.
- Establish tagging conventions (by application, business unit, or environment) so tasks, workloads, and storage assets can be tracked consistently across teams.

Who to involve:

- **Storage administrator:** Sets up the ticketing system as the intake system of record, deploys Pure1 API integrations, and configures tagging conventions across all storage assets
- **Infrastructure architect:** Defines the OpenMetrics telemetry schema, establishes the centralized dashboard architecture, and sets the tagging taxonomy to ensure consistent asset tracking across teams
- **IT operations lead:** Owns the operational tooling stack selection, ensures the collaboration workspace is stood up and adopted, and connects Pure1 telemetry to the monitoring platform

Success indicators:

- 100% of infrastructure work requests originate in the ticketing system.
- A single dashboard view shows open, in-progress, and completed ops tasks across the storage estate.
- Runbook content is centralized and versioned; no critical procedures live only in individual inboxes or local drives.
- Decisions can be made with fleet-wide, contextual metrics.

Everpure enablers:

- OpenMetrics exporter embedded in Purity to send detailed telemetry to the monitoring platform
- Pure1 telemetry and APIs to run predictive forecasting and generate tickets based on expected capacity or performance issues
- Tag-based workload identification in Purity and arrays in Pure1 so assets and workloads can easily be identified and tracked

Process maturity

Requirement: Introduce project tracking and functional accountability for delivery activities.

Recommended actions:

- Define a lightweight intake-to-delivery workflow with named stages (request → triage → assigned → in progress → validated → closed), and apply it consistently to all provisioning, change, and decommission work.
- Assign a single named owner to every active work item; eliminate “the team will handle it” as an acceptable assignee.
- Introduce a weekly delivery standup where active tickets are reviewed against stage and owner, replacing reactive hallway coordination.
- Document the top 10 most frequent operational tasks (volume provisioning, host attach, snapshot scheduling, replication setup, and so on) as repeatable procedures with defined inputs, steps, and exit criteria.
- Capture cycle time (intake to closure) for each work item type so baseline delivery performance is visible before optimization begins.

Who to involve:

- **Storage administrator:** Documents the top 10 recurring operational procedures, implements the intake-to-delivery workflow, and ensures every active work item carries a named owner and current stage
- **IT operations lead:** Introduces and chairs the weekly delivery standup, establishes the cycle time baseline for recurring task types, and drives consistent use of documented procedures
- **Infrastructure architect:** Reviews the documented procedures for accuracy and completeness, validates that Everpure Fusion presets map to the defined repeatable tasks, and confirms that workload simulation outputs are used for placement decisions

Success indicators:

- Every open work item has a named owner and a current stage at any point in time.
- The top 10 recurring tasks have documented procedures referenced in tickets rather than reinvented per request.
- Baseline cycle times exist for at least three recurring task types within 90 days.

Everpure enablers:

- Everpure Fusion presets for repeatable provisioning tasks rather than using manual processes or writing ad hoc scripts
- Pure1 workload simulations to run analytics on workloads prior to deployment for documented placement decisions
- Tag-based asset and workload categorization to group similar requests and ensure delivery, auditing, and reporting to proper end users and departments

Governance and measurement

Requirement: Assign execution leads per function; build recurring ops reviews into team cadences.

Recommended actions:

- Designate an execution lead for each major operational function (provisioning, data protection, performance, capacity) with explicit accountability for backlog health and delivery against SLAs in that area.
- Establish a monthly operations review where execution leads present critical open work, recent incidents, capacity trajectory, and outstanding risks to infrastructure leadership.
- Define a small starter set of operational KPIs (open ticket aging, mean time to provision, change success rate, capacity headroom) and make them standing agenda items in the review.
- Build a simple responsible, accountable, consulted, and informed (RACI) matrix for the most common operational scenarios so escalation paths and decision rights are unambiguous.
- Begin tracking which incidents and delays had downstream business impact, even informally, to seed the business-awareness shift expected at higher maturity levels.

Who to involve:

- **Head of Infrastructure:** Designates execution leads per operational function, establishes the monthly operations review cadence, and ensures escalation paths and decision rights are documented in the RACI matrix
- **Execution leads (per function):** Own backlog health and SLA delivery for their area, present findings at the monthly ops review, and start tracking incidents with downstream business impact
- **IT service manager:** Maintains the RACI matrix, manages the ops review agenda, and ensures operational KPIs are reported on cadence with trend data

Success indicators:

- Every operational function has a named execution lead known across the team.
- A recurring ops review happens on a defined cadence with consistent attendance and a published agenda.
- At least four operational KPIs are reported on every cycle with trend data over time.

Everpure enablers:

- Everpure Fusion fleet-wide reporting for cross-array KPIs
- Pure1 historical telemetry for trend and risk analysis
- Tag-based grouping to slice metrics by function, business unit, or environment

Benefits and metrics

Requirement: Execution is often delayed or reactive; visibility is limited to siloed workflows.

Recommended actions:

- Establish baseline measurements for mean time to provision, mean time to resolve (MTTR), and ticket aging within the first 30 days so improvement can be demonstrated against a real starting point.
- Build a single operational scorecard that surfaces delivery throughput, backlog size, and SLA adherence in one view.
- Track the proportion of work that is planned versus reactive each month to make the shift away from incident-driven operations visible.
- Capture qualitative feedback from at least one internal application owner per quarter on storage delivery experience, providing an early signal that complements the quantitative metrics.
- Publish the scorecard to infrastructure leadership and the execution leads so metrics are visible, not buried.

Who to involve:

- **IT operations lead:** Establishes the baseline scorecard within the first 30 days, tracks planned versus reactive work ratio monthly, and ensures metrics are published to leadership and execution leads
- **Storage administrator:** Provides raw data on mean time to provision, MTTR, and ticket aging to feed the operational scorecard
- **Application owner:** Provides quarterly qualitative feedback on storage delivery experience, giving an early signal that complements quantitative metrics

Success indicators:

- A baseline exists for at least three delivery metrics by day 30 and is refreshed monthly thereafter.
- The planned versus reactive work ratio is reported and trending in a healthier direction over two consecutive quarters.
- Visibility into delivery performance is no longer confined to individual team leads.

Everpure enablers:

- OpenMetrics exporter to gather relevant data across all storage assets to feed into operational dashboards as the primary scorecard surface
- Tag conventions enabling per-application or per-business unit slicing and reporting

Everpure features

Requirement: Everpure Fusion operational dashboards, project tracking via tags, and Pure1 delivery status telemetry

Recommended actions:

- Enable Everpure Fusion across the storage estate and configure OpenTelemetry-based operational dashboards as the primary day-to-day view for the ops team, replacing per-array consoles for routine visibility.
- Define and build your Everpure Fusion fleet to resemble the physical and logical layouts of your data centers and public clouds for easier workload placement and planning.
- Define and roll out a tagging taxonomy covering, at minimum, application, business unit, environment (production/nonproduction), and criticality tier; apply it to existing workloads and storage arrays as well as new provisioning.
- Connect your internal ticketing system with Pure1 via APIs so operational, capacity, and performance events feed the ticketing system and the scorecard, closing the loop between array and workload state and tracked work.
- Configure role-based access so execution leads, operators, and leadership all see views appropriate to their accountability rather than a single shared dashboard.
- Document which Pure1 and Everpure Fusion capabilities map to which operational outcomes (this becomes the foundation for the more advanced automation expected in the ML-2 → ML-3 transition).

Who to involve:

- **Storage administrator:** Enables Everpure Fusion across the estate, configures OpenTelemetry dashboards as the primary ops view, and rolls out the tagging taxonomy to existing workloads and new provisioning
- **Infrastructure architect:** Defines Everpure Fusion fleet topology to mirror the physical and logical data center layout, sets role-based access controls for execution leads and operators, and maps Pure1 and Everpure Fusion capabilities to operational outcomes
- **IT operations lead:** Connects the internal ticketing system to Pure1 via API, validates that the array and workload state feed both the ticketing system and the operational scorecard, and confirms the feedback loop is closed

Success indicators:

- OpenTelemetry-based dashboards are the team's default landing surface for daily operations and tasks.
- Tagging coverage exceeds an agreed threshold (for example, 90% of provisioned capacity carries the full taxonomy) within 90 days.
- Pure1 telemetry is demonstrably wired into both the ticketing system and the operational scorecard.

Everpure enablers:

- Everpure Fusion operational dashboards
- Pure1 telemetry and API integration with ticketing systems
- Pure1 and Purity tagging for workload and cost attribution
- Role-based access controls in Everpure Fusion

Maturity path: Level 2 → Level 3

The move to automated deployment

Transitioning to ML-3 operationalizes application deployment by introducing automated infrastructure builds for mission-critical applications, aligning them directly to expected SLAs and business outcomes.

Category	Requirement
Technology enhancements	Enable workflow automation for approvals, resource requests, and deployment tasks.
Process maturity	Standardize delivery runbooks; define metrics for execution predictability and team throughput.
Governance and measurement	Operations and delivery managers coordinate across departments; SLAs are embedded into execution KPIs.
Benefits and metrics	Throughput and backlog are tracked; delivery begins to align with business timelines.
Everpure features	Everpure Fusion automation templates, task telemetry in Pure1, and cross-team coordination controls

Technology enhancements

Requirement: Enable workflow automation for approvals, resource requests, and deployment tasks.

Recommended actions:

- Implement automated approval workflows in the ticketing platform so standard requests (volume provisioning under a defined size, snapshot policy changes, host group additions) are routed, approved, and executed without manual intervention.
- Adopt Everpure Fusion presets for the most common provisioning patterns identified during ML-2, converting documented runbooks into executable automation.
- Integrate the ticketing system with Everpure Fusion via API so approved tickets trigger provisioning actions directly, eliminating the manual handoff between request and execution.
- Stand up infrastructure-as-code patterns (Ansible, Terraform, or PowerShell modules) for repeatable storage configurations, version-controlled in a shared repository.
- Enable self-service request portals for application owners to submit standard resource requests against preapproved service catalogs, with guardrails defined by policy rather than reviewed case by case.

Who to involve:

- **Storage administrator:** Implements automated approval workflows in the ticketing platform, adopts Everpure Fusion presets for high-volume provisioning patterns, and integrates the ticketing system with Everpure Fusion via API
- **Automation engineer:** Builds and maintains the infrastructure-as-code patterns in a shared repository, configures the self-service request portal, and validates that approved tickets trigger provisioning actions without manual handoff
- **Infrastructure architect:** Defines guardrails and service catalog structure for the self-service portal, reviews Everpure Fusion automation templates before production promotion, and ensures infrastructure-as-code patterns align with the overall infrastructure design

Success indicators:

- At least 60% of standard provisioning requests complete without manual touch within 90 days.
- The mean time to provision for standard requests drops by 50% or more against the ML-2 baseline.
- All Everpure Fusion automation templates are version-controlled and reviewed before promotion to production use.

Everpure enablers:

- Everpure Fusion presets
- Everpure Fusion REST API for ticketing system integration
- Pure1 audit logs to confirm automated execution outcomes

Process maturity

Requirement: Standardize delivery runbooks; define metrics for execution predictability and team throughput.

Recommended actions:

- Convert the top 10 procedures documented during ML-2 into standardized runbooks with defined SLAs, owners, escalation paths, and validation steps; these become the contract between storage ops and application teams.
- Establish a runbook review cadence (quarterly) to retire stale procedures, incorporate lessons from incidents, and promote ad hoc procedures that have proven repeatable.
- Define throughput metrics beyond cycle time: automation coverage percentage, exception rate (requests requiring manual intervention), and rework rate.
- Introduce a deployment readiness checklist for mission-critical applications that includes capacity validation, replication configuration, data protection alignment, and SLA confirmation before go-live.
- Adopt a change advisory process for nonstandard work that distinguishes routine automated changes from those requiring human review, preventing the change advisory board from becoming a bottleneck for low-risk work.

Who to involve:

- **Storage administrator:** Converts the top 10 documented procedures into versioned runbooks with defined SLAs, owners, escalation paths, and validation steps, and maintains the quarterly runbook review cadence
- **IT service manager:** Manages the change advisory process, distinguishes routine automated changes from those requiring human review, and tracks throughput metrics including exception rate and rework rate
- **Automation engineer:** Measures automation coverage percentage and exception rates, and ensures deployment readiness checklists are completed for mission-critical applications before go-live

Success indicators:

- 100% of routine operational procedures execute against a published, versioned runbook.
- Execution predictability metrics (estimated versus actual delivery time) are reported, with variance trending downward over two consecutive quarters.
- The exception rate for standard requests stays below an agreed threshold (for example, 15%).

Everpure enablers:

- Everpure Fusion presets as the executable form of standardized runbooks
- Pure1 audit logs to measure predictability and throughput
- Everpure Fusion fleet topology groups for consistent runbook execution across regions, availability zones/data centers, and subgroups

Governance and measurement

Requirement: Operations and delivery managers coordinate across departments; SLAs are embedded into execution KPIs.

Recommended actions:

- Establish formal coordination forums between storage operations and adjacent functions (compute, network, application delivery, security) on a defined cadence, replacing the ad hoc cross-team interactions that worked at a smaller scale.
- Define per-application or per-tier SLAs for provisioning time, availability, RPO, and RTO, with explicit ownership and consequences for breach.
- Embed SLA adherence into the operational scorecard introduced at ML-2, with red/yellow/green status visible to leadership and execution leads.
- Implement a shared change calendar across infrastructure functions so storage, network, and compute changes are coordinated rather than colliding in production windows.
- Define and document service tiers (for example, Minimum Viable Business, Tier 1/2/3, Mission-Critical, Business-Important, Standard) with differentiated SLAs, response times, and automation coverage so investment matches business value.

Who to involve:

- **Head of Infrastructure:** Chairs cross-functional coordination forums with storage, compute, network, application delivery, and security, and maintains the shared change calendar across infrastructure domains
- **Application owner/LoB lead:** Validates per-application and per-tier SLAs for provisioning time, availability, RPO, and RTO, and reviews SLA adherence reporting to hold infrastructure accountable to business commitments
- **IT service manager:** Embeds SLA adherence into the operational scorecard with red/yellow/green status, documents service tier definitions, and ensures the differentiated SLA structure matches business value

Success indicators:

- Cross-functional coordination forums meet on cadence, with documented decisions and action items.
- SLA adherence is reported per service tier and per application with at least 90 days of trend data.
- Change collisions across infrastructure domains drop measurably against the ML-2 baseline.

Everpure enablers:

- Everpure Fusion presets aligned to application and service tier SLAs
- Pure1 and Purity OpenMetrics telemetry to track SLA breaches
- Tag-based service tier and application identification on arrays, workloads, and volumes

Benefits and metrics

Requirement: Throughput and backlog are tracked; delivery begins to align with business timelines.

Recommended actions:

- Expand the operational scorecard to show delivery against business-aligned milestones (quarterly project commitments, application release cycles) rather than only internal ops metrics.
- Track automation coverage as a headline metric—percentage of work executed via Everpure Fusion presets and API rather than manual procedure—and set a quarterly target.
- Measure the mean time to provision against tiered SLAs, with separate reporting for mission-critical workloads and standard tiers.
- Capture business impact data alongside operational metrics, including number of application releases enabled, time-to-market improvement for new workloads, and infrastructure-related deployment delays avoided.
- Introduce a quarterly review with application and business stakeholders to validate that storage delivery is matching business cadence, not just ops cadence.

Who to involve:

- **IT operations lead:** Expands the operational scorecard to show delivery against business-aligned milestones, tracks automation coverage as a headline metric, and runs quarterly reviews with application and business stakeholders
- **Storage administrator:** Provides mean time to provision data segmented by service tier and tracks backlog aging against SLA commitments
- **Application owner/LoB lead:** Supplies business impact data alongside operational metrics, including application releases enabled and infrastructure-related deployment delays avoided

Success indicators:

- Automation coverage rises consistently quarter over quarter toward a defined ML-3 target (for example, 70%).
- Backlog aging shows that work is completed within the SLA for the relevant tier rather than aging indefinitely.
- At least one business impact metric (release enablement, time to market, and so on) is reported alongside operational throughput each cycle.

Everpure enablers:

- Operational dashboards extended with automation coverage and SLA views by measuring work done via Everpure Fusion presets and APIs
- Pure1 telemetry and audit logs segmented by tag and service tier

Everpure features

Requirement: Everpure Fusion automation templates, telemetry in Pure1, and Pure1 and Purity tags

Recommended actions:

- Build a library of Everpure Fusion presets covering the highest-volume provisioning patterns, with presets promoted into production through a defined review and testing process.
- Monitor Pure1 telemetry and audit logs via API to flow into the operational scorecard and the ticketing system, providing closed-loop confirmation that automated tasks completed successfully and within expected duration.
- Establish an Everpure Fusion preset catalog tied to service tiers, so requesting a “Mission-Critical Tier 1” volume applies the right protection group, replication policy, performance configuration, and tagging automatically.
- Integrate Pure1 AI-driven recommendations (workload placement, capacity forecasting, performance optimization) into the operational review cadence, surfacing them as outputs to capacity and performance planning.

Who to involve:

- **Storage administrator:** Builds the Everpure Fusion preset library covering the highest-volume provisioning patterns, promotes presets through a defined review process, and establishes the service tier-aligned preset catalog
- **Automation engineer:** Wires Pure1 telemetry and audit logs via API into the ticketing system and operational scorecard, and integrates Pure1 AI-driven recommendations into the operational review cadence
- **Infrastructure architect:** Defines Everpure Fusion fleet topology groups for consistent preset execution across regions and data centers, and ensures all arrays, workloads, and volumes carry the full tagging taxonomy

Success indicators:

- The Everpure Fusion preset library covers the patterns responsible for at least 70% of provisioning volume.
- Pure1 telemetry is consumed by both the ticketing system and the operational scorecard via API, with no manual export steps.
- All arrays are tagged in Pure1 and all workloads and volumes are tagged in Purity for tracking and monitoring.

Everpure enablers:

- Everpure Fusion preset library and service tier catalog
- Pure1 telemetry and audit logs via API
- Pure1 AI-driven recommendations for workload placement and capacity forecasting
- Everpure Fusion fleet topology groups for consistent preset execution

Maturity path: Level 3 → Level 4

The move to orchestrated application deployment

At ML-4, automation becomes workload-aware, with end-to-end orchestration integrated into real-time monitoring and business change cycles.

Category	Requirement
Technology enhancements	Deploy adaptive planning platforms, resource simulation tools, and execution telemetry.
Process maturity	Align quarterly planning with delivery cycles; integrate readiness metrics into OKRs.
Governance and measurement	An execution readiness committee oversees cross-functional delivery; risk is actively tracked and managed.
Benefits and metrics	Backlogs are cleared predictably; pivot speed improves with integrated execution signals.
Everpure features	Everpure Fusion project forecasting, Pure1 execution velocity analytics, and readiness scoring modules

Technology enhancements

Requirement: Deploy adaptive planning platforms, resource simulation tools, and execution telemetry.

Recommended actions:

- Deploy and configure new storage assets automatically following established array templates to meet assigned tier and application SLAs and business outcomes.
- Implement workload-aware orchestration that ties infrastructure provisioning decisions to application context; performance profiles, data protection requirements, and business criticality drive placement and configuration automatically rather than being specified per request.
- Deploy resource simulation capabilities (Pure1 workload simulator, capacity forecasting models) to model the impact of proposed changes before they execute, replacing trial-and-error provisioning with informed placement.
- Integrate telemetry and audit logs from Pure1 and Purity with broader observability platforms (Datadog, Splunk, ServiceNow ITOM) so storage signals participate in the same operational picture as compute, network, and application telemetry.
- Adopt closed-loop automation patterns where telemetry triggers remediation: predicted capacity exhaustion triggers procurement workflows, performance degradation triggers workload rebalancing, and replication lag triggers operator notification with suggested actions.
- Establish a CI/CD pipeline for infrastructure changes with Everpure Fusion presets and infrastructure-as-code modules promoted through dev/test/prod environments rather than edited in place.

Who to involve:

- **Infrastructure architect:** Implements workload-aware orchestration that ties provisioning to application context, deploys resource simulation capabilities, and establishes the CI/CD pipeline for infrastructure changes
- **Automation engineer:** Builds and maintains closed-loop automation patterns, integrates Pure1 and Purity telemetry with third-party observability platforms, and ensures Purity zero-touch provisioning is active for new array deployment
- **Site reliability engineer:** Configures storage telemetry to participate in cross-domain observability views used by infrastructure, application, and business stakeholders, and validates that closed-loop automation handles at least three scenarios without human intervention
- **Data scientist/ML engineer:** Supports integration of Pure1 workload simulation outputs and capacity forecasting models to drive simulated placement decisions for new deployments

Success indicators:

- Workload placement decisions are driven by simulated outcomes rather than rule of thumb or operator preference for at least 80% of new deployments.
- Storage telemetry is consumed in cross-domain observability views used by infrastructure, application, and business stakeholders.
- Closed-loop automation handles at least three operational scenarios end to end without human intervention.
- New storage arrays are ready for consumption in minutes and can be audited against established templates based on purpose and assigned tier.

Everpure enablers:

- Pure1 workload simulator and capacity forecasting
- Everpure Fusion adaptive policy engine for workload-aware placement and rebalancing
- Pure1 telemetry export and integration with third-party observability platforms
- Purity zero-touch provisioning for new array deployment

Process maturity

Requirement: Align quarterly planning with delivery cycles; integrate readiness metrics into OKRs.

Recommended actions:

- Integrate infrastructure planning cycles with application and business planning cycles so capacity, performance, and data protection commitments are negotiated as part of the same quarterly cadence rather than separately.
- Define readiness metrics that quantify infrastructure's ability to meet upcoming business demand: forecasted capacity headroom against planned workload growth, performance margin against expected load, and recoverability against tested scenarios.
- Embed these readiness metrics in team OKRs at the infrastructure leadership level, with measurable quarterly targets tied to business outcomes.
- Establish a quarterly business review (QBR) cadence between infrastructure and major application or business unit stakeholders, with shared readiness dashboards as the basis for the conversation.
- Adopt scenario-based planning practices: model the impact of major upcoming events (product launches, mergers and acquisitions, fiscal year close) on infrastructure capacity and performance, then provision proactively.

Who to involve:

- **Head of Infrastructure:** Integrates infrastructure planning into the quarterly business cadence, embeds readiness metrics into team OKRs, and establishes the QBR with major application and business unit stakeholders
- **Infrastructure architect:** Leads scenario-based planning for major business events, validates capacity and performance commitments against planned workload growth, and documents preprovisioning decisions with post-event reviews
- **IT service manager:** Manages the shared readiness dashboards used in QBRs and ensures delivery cycle commitments are tracked against specific business initiatives rather than generic capacity additions

Success indicators:

- Quarterly planning produces infrastructure commitments tied to specific business initiatives, not generic capacity additions.
- Readiness metrics appear in team and leadership OKRs and are reported against quarterly metrics.
- At least two major business events per year are premodeled and preprovisioned for, with post-event reviews validating the model.

Everpure enablers:

- Everpure Fusion preset reviews for forecasting tied to planning cycles
- Pure1 capacity and performance forecasting for scenario modeling
- Readiness scoring modules exposed in dashboards consumed by leadership

Governance and measurement

Requirement: An execution readiness committee oversees cross-functional delivery; risk is actively tracked and managed.

Recommended actions:

- Establish a formal execution readiness committee with representation from storage, compute, network, application delivery, security, and a business sponsor; charter it to oversee cross-functional delivery commitments and resolve cross-domain risks.
- Implement an active risk register for infrastructure delivery, with risks categorized by likelihood, impact, and time to mitigate; review the register at every committee meeting, with decisions and owners captured.
- Define risk thresholds that trigger automated escalation (for example, capacity headroom below tier threshold for more than 30 days escalates to committee or SLA breach in mission-critical tier escalates immediately).
- Adopt a portfolio view of in-flight delivery work so the committee can prioritize and rebalance across initiatives rather than evaluating requests in isolation.
- Implement post-incident and post-delivery reviews that feed lessons learned into runbooks, automation templates, and risk register entries, closing the loop between operational reality and governance.

Who to involve:

- **Head of Infrastructure:** Chairs the execution readiness committee, maintains the active risk register, and ensures cross-functional dependencies are visible in the portfolio view used for trade-off decisions
- **Enterprise application architect:** Represents application delivery in the committee, escalates cross-domain risks, and ensures post-delivery reviews feed lessons learned back into runbooks and automation templates
- **CIO/CTO:** Provides executive sponsorship for the execution readiness committee, validates risk thresholds that trigger automated escalation, and reviews the portfolio view for strategic prioritization
- **Security lead:** Ensures security requirements are represented in cross-functional delivery commitments and that risks with security implications are properly categorized and mitigated in the risk register

Success indicators:

- The execution readiness committee meets on a defined cadence, with documented decisions, action items, and risk register updates.
- The risk register is actively maintained, with measurable reduction in aged, unmitigated risks over time.
- Cross-functional dependencies are visible in a portfolio view that the committee uses to make trade-off decisions.

Everpure enablers:

- Purity and Everpure Fusion fleet-wide risk and readiness reporting
- Pure1 replication, capacity, and performance trend data feeding the risk register
- Readiness scoring modules providing committee-level visibility into delivery health

Benefits and metrics

Requirement: Backlogs are cleared predictably; pivot speed improves with integrated execution signals.

Recommended actions:

- Measure pivot speed explicitly—the time to replan and redeploy infrastructure in response to a business change (acquisition, divestiture, market shift, application rearchitecture), with documented before and after examples.
- Report backlog clearance against forecast—not just whether work was completed, but whether completion matched the rate projected at the start of the cycle, validating predictability.
- Introduce leading indicators alongside lagging metrics: forecasted SLA breaches, projected capacity exhaustion dates, and predicted performance bottlenecks. The presence of leading indicators is itself an ML-4 marker.
- Track decision latency—the time from a business question being asked (“Can we onboard this acquired company in 60 days?”) to a defensible infrastructure answer being delivered—and drive it down with simulation and forecasting.
- Quantify business value enabled rather than only cost incurred: revenue protected by SLA adherence, risk averted, time-to-market acceleration, deployment frequency increase, and infrastructure-related deployment failures avoided.

Who to involve:

- **IT operations lead:** Measures pivot speed for real business events, reports backlog clearance against forecast, and tracks decision latency from business question to defensible infrastructure answer
- **Financial analyst:** Quantifies business value enabled rather than only cost incurred, including revenue protected by SLA adherence, risk averted, and time-to-market acceleration
- **Site reliability engineer:** Introduces and reports leading indicators alongside lagging metrics, including forecasted SLA breaches, projected capacity exhaustion dates, and predicted performance bottlenecks

Success indicators:

- Pivot speed is documented for at least three real business events, with measurable improvement over the ML-3 baseline.
- Leading indicators are reported alongside lagging metrics in every operational review.
- Business value metrics (not just operational metrics) appear in infrastructure leadership scorecards.

Everpure enablers:

- Pure1 forecast planning and analytics to formulate predictive readiness signals
- Readiness scoring modules combining capacity, performance, and protection signals into a single composite view

Everpure features

Requirement: Everpure Fusion project forecasting, Pure1 execution velocity analytics, and readiness scoring modules

Recommended actions:

- Enable Pure1 forecasting to model in-flight and planned work against fleet capacity, performance, and protection commitments, with forecasts refreshed on a defined cadence and surfaced to the execution readiness committee.
- Use Pure1 analytics and audit logs to track delivery throughput trends, identify slowdowns, and correlate operational performance with upstream changes (new array deployments, automation template updates, organizational changes).
- Use zero-touch provisioning to rapidly deploy new storage assets against array templates based on array purpose and tier, and track mean time to array delivery to production readiness.
- Deploy readiness scoring modules to produce composite scores per application, service tier, or business unit; these single numbers summarize whether storage is ready to meet upcoming demand and are decomposable into capacity, performance, protection, and operational dimensions.
- Integrate Pure1 AI workload recommendations into orchestration decisions: when an application owner requests a new workload, Everpure Fusion consumes the Pure1 placement recommendation as the default, with operator override available but tracked and audited.
- Establish a feedback loop where forecast accuracy is measured and improved over time, treating the forecasting platform itself as a system that needs tuning.

Who to involve:

- **Infrastructure architect:** Enables Pure1 forecasting for in-flight and planned work, uses zero-touch provisioning for new array deployment against templates, and tracks mean time to array delivery to production readiness
- **Storage administrator:** Deploys readiness scoring modules per application, service tier, and business unit, and ensures Pure1 AI workload recommendations are consumed by Everpure Fusion as the default for new workload placement
- **IT operations lead:** Ensures forecast accuracy is measured and improved over time, validates that readiness scores are visible to application owners and business stakeholders, and reviews Pure1 analytics in every execution readiness committee meeting

Success indicators:

- Pure1 and project forecasts are reviewed in every execution readiness committee meeting and inform at least one major decision per quarter.
- Readiness scores are visible to application owners and business stakeholders, not just infrastructure staff.
- Forecast accuracy (predicted versus actual) is measured and trending upward over consecutive quarters.
- The majority (70%+) of new workloads are deployed to storage assets automatically via presets and Everpure Fusion leveraging Pure1 workload simulations.

Everpure enablers:

- Pure1 forecasting and execution velocity analytics
- Readiness scoring modules for application, service tier, and business unit views
- Pure1 AI workload recommendations integrated into Everpure Fusion orchestration
- Purity zero-touch provisioning for new array deployment

Maturity path: Level 4 → Level 5

The move to autonomous, self-healing resilience

ML-5 is the pinnacle state, where resilience is self-healing, systems autonomously reroute services during disruption, and recovery metrics are fully transparent to leadership.

Category	Requirement
Technology enhancements	Adopt AI-based execution modeling, cross-pillar orchestration, and proactive risk avoidance systems.
Process maturity	Use predictive readiness and task dependency engines to plan and deliver adaptively.
Governance and measurement	There is board-level insight into execution constraints; scenario planning is embedded into governance cycles.
Benefits and metrics	Near-real-time agility is achieved; dependencies are resolved automatically ahead of bottlenecks.
Everpure features	Everpure Fusion readiness engine, Pure1 adaptive risk modeling, and intelligent orchestration tools

Technology enhancements

Requirement: Adopt AI-based execution modeling, cross-pillar orchestration, and proactive risk avoidance systems.

Recommended actions:

- Deploy AI-based execution modeling that learns from historical operational data—provisioning outcomes, performance anomalies, capacity trajectories, and incident patterns—and produces probabilistic forecasts that improve over time without manual retuning.
- Implement cross-pillar orchestration that coordinates infrastructure changes across storage, compute, network, security, and data protection in a single workflow rather than separate domain-specific pipelines, with dependencies resolved automatically.
- Adopt proactive risk avoidance with systems that detect emerging conditions (drift toward capacity exhaustion, replication lag patterns, performance degradation signatures) and trigger remediation before service impact occurs, with operator notification rather than operator action.
- Establish a policy-as-code framework where business intent (SLAs, compliance requirements, cost constraints, data residency rules) is expressed in code and enforced automatically across the fleet, with violations producing both alerts and auto-remediation where safe.
- Build a digital twin of the infrastructure environment that allows changes to be modeled, validated, and stress-tested in simulation before being applied to production, with the twin kept current automatically from live telemetry.

Who to involve:

- **Enterprise application architect:** Deploys the policy-as-code framework expressing business intent across the fleet and defines the digital twin environment for preproduction change modeling
- **Data scientist/ML engineer:** Deploys and maintains the AI-based execution modeling system, ensures models improve over time without manual retuning, and monitors model accuracy and drift
- **Site reliability engineer:** Validates proactive risk avoidance by tracking incidents prevented per quarter, audits cross-pillar orchestration workflows, and confirms that policy enforcement produces audit trails for every action
- **Head of Infrastructure/CIO:** Sets the boundaries of autonomous operation, approves the scope of declarative policy-as-code coverage, and ensures the digital twin is kept current from live telemetry

Success indicators:

- AI-driven recommendations are accepted and acted upon (automatically or with operator approval) for the majority of orchestration decisions.
- Proactive risk avoidance prevents measurable incidents per quarter that would otherwise have caused service impact, with documented examples.
- Policy-as-code coverage spans capacity, performance, protection, security, and compliance, with audit trails for every policy enforcement and exception.

Everpure enablers:

- Pure1 adaptive risk modeling and AI-driven recommendations
- Everpure Fusion policy engine for declarative, code-driven intent
- Purity APIs leveraged extensively to treat infrastructure as code

Process maturity

Requirement: Use predictive readiness and task dependency engines to plan and deliver adaptively.

Recommended actions:

- Adopt continuous planning rather than quarterly planning as the default cadence; readiness signals, business inputs, and dependency state continuously update the delivery plan, with humans intervening to set direction and resolve ambiguity rather than to schedule work.
- Deploy task dependency engines that map relationships between in-flight work items, planned changes, and business commitments, exposing critical paths and at-risk dependencies in real time.
- Establish autonomous delivery for well-understood work categories; requests that fit known patterns, fall within policy guardrails, and have acceptable risk profiles execute end to end without human approval, with humans engaged only for exceptions.
- Implement chaos engineering practices for storage and data services, including deliberate, controlled fault injection that validates self-healing behaviors and surfaces resilience gaps before real incidents do.
- Treat the operating model itself as a system under management, with metrics on how often human intervention was required, where it was required, and why, that feed back into automation coverage and policy refinement.

Who to involve:

- **Site reliability engineer:** Leads the chaos engineering discipline, executes periodic fault injection tests to validate self-healing behaviors, and ensures findings feed back into automation coverage and policy refinement
- **Storage architect/storage administrator:** Defines intent-based policies and service catalog structures that the autonomous remediation modules use, and reviews autonomy boundaries on a defined cadence
- **Financial analyst/CFO:** Designs, implements, and governs the full chargeback model, defines cost-to-serve and unit economics calculations, and validates that storage is tracked as a financial service at the business unit level

Success indicators:

- The proportion of work delivered fully autonomously (no human approval or manual steps) exceeds an agreed threshold for in-scope categories.
- Task dependency engines surface at least one critical path adjustment per cycle that would have been missed by manual planning.
- Chaos engineering exercises run on a defined cadence, with documented findings feeding back into automation and policy improvements.

Everpure enablers:

- Pure1 telemetry and workload simulations for adaptive risk modeling feeding dependency analysis
- Intelligent orchestration tools executing autonomous delivery within policy bounds using Everpure Fusion presets
- Placement and rebalancing decisions driven by Everpure Fusion using Pure1 workload simulations and forecasting

Governance and measurement

Requirement: There is board-level insight into execution constraints; scenario planning is embedded into governance cycles.

Recommended actions:

- Establish a board-level operational view that presents infrastructure execution constraints in business terms—revenue at risk from capacity shortfalls, time-to-market exposure from performance bottlenecks, and compliance posture across regions—refreshed continuously rather than reported quarterly.
- Embed scenario planning into the governance cadence at the highest level: every major board-relevant decision (mergers and acquisitions, geographic expansion, product launch, divestiture) is modeled against the infrastructure baseline before commitment, with quantified readiness and risk outputs.
- Implement explicit accountability for autonomous decisions: every action taken by the system is logged, attributable, reviewable, and tied back to the policy or model that authorized it, supporting both audit and continuous improvement.
- Define and govern the boundaries of autonomy: which categories of action the system may take without approval, under which conditions human review is mandatory, and how those boundaries are reviewed and adjusted over time.
- Adopt a continuous assurance model where compliance, resilience, and SLA posture are validated continuously by the system itself, with attestation artifacts produced automatically rather than assembled for audits.

Who to involve:

- **CIO/CTO:** Establishes the board-level operational view in business terms, embeds storage efficiency and availability OKRs into business unit performance, and monitors cost-to-serve and forecast variance
- **Application owner/LoB lead:** Holds accountability for efficiency and availability OKRs embedded in business unit's performance metrics and drives consumption alignment with business needs
- **Financial analyst:** Compiles and validates financial metrics including cost-to-serve, unit economics, and forecast variance, and presents them to executive leadership for strategic decision-making
- **Head of Infrastructure:** Ensures continuous assurance compliance posture is validated by the system itself and that attestation artifacts are produced automatically rather than assembled for audits

Success indicators:

- A board-level operational view exists and is consulted in at least one major strategic decision per quarter.
- Scenario planning outputs accompany major business commitments, with documented infrastructure readiness assessments.
- Autonomy boundaries are formally defined, reviewed on cadence, and demonstrably enforced, with full audit trail coverage.

Everpure enablers:

- Purity OpenMetrics exporter and APIs provide telemetry to feed into executive views.
- Pure1 forecasting, historical data, and trending via API feed into adaptive risk modeling to support scenario analysis.
- Audit capabilities across Purity and Pure1 provide autonomous action accountability.
- The Everpure Fusion compliance engine ensures workloads and assets remain aligned to SLAs and business outcomes.

Benefits and metrics

Requirement: Near-real-time agility is achieved; dependencies are resolved automatically ahead of bottlenecks.

Recommended actions:

- Measure time to decision and time to execution for business-driven infrastructure changes in hours or minutes rather than days, and report these as headline metrics alongside availability and performance.
- Track autonomous resolution rate: the percentage of operational conditions (capacity pressure, performance anomalies, protection gaps, dependency conflicts) resolved without human action, with the goal of steady improvement quarter over quarter.
- Quantify avoided incidents and delays—outcomes that did not happen because proactive risk avoidance intervened—using counterfactual analysis where possible.
- Report agility in business terms: deployment frequency, lead time for new applications or environments, mean time to recover from major disruption, and percentage of business-driven changes delivered within the requesting cycle.
- Continuously measure model and forecast quality—accuracy, calibration, and drift— and treat improvements to model quality as first-class operational outcomes.

Who to involve:

- **Financial analyst/CFO:** Quantifies cost-to-serve and unit economics across business units, minimizes forecasting variance through accurate financial modeling, and reports avoided-incident value using counterfactual analysis
- **Site reliability engineer:** Tracks and reports the autonomous resolution rate, measures time to decision and time to execution for business-driven infrastructure changes, and validates that deployment frequency and lead time are reported in business terms
- **CIO/CTO:** Ensures headline operational metrics are accessible to business stakeholders in near-real time, not gated behind monthly cycles, and champions continuous model quality measurement as a first-class operational outcome

Success indicators:

- Headline operational metrics are reported in near-real time and accessible to business stakeholders, not gated behind monthly cycles.
- The autonomous resolution rate is measured, reported, and trending upward against a defined baseline.
- Avoided-incident reporting is in place with credible methodology, providing a quantified view of proactive risk avoidance value.

Everpure enablers:

- Pure1 intelligent AI forecasting for adaptive risk analytics
- Purity OpenMetrics exporter and APIs for cross-fleet searchability and observability and continuous real-time metrics
- Pure1 audit and orchestration telemetry for autonomous action measurement

Everpure features

Requirement: Everpure Fusion compliance engine, Pure1 adaptive risk modeling, and intelligent orchestration tools

Recommended actions:

- Operate the Everpure Fusion compliance engine as the authoritative source for fleet-wide workload state, with continuous signals feeding planning, governance, executive reporting, and autonomous delivery decisions.
- Configure the Everpure Fusion compliance engine to automatically alert on SLA drift on all workloads across the entire fleet, prompting user or automated interaction to remediate drift.
- Deploy Pure1 regular, automated forecasting as the core analytics layer feeding into agentic AI tools that learn from historical outcomes, anticipate emerging risk conditions, and produce actionable recommendations consumed by both humans and the orchestration platform.
- Operationalize intelligent orchestration tools as the execution layer: workload-aware, policy-enforced, and cross-pillar, with autonomous action for in-scope categories and human-in-the-loop escalation for exceptions.
- Integrate Pure1 and Everpure Fusion capabilities with adjacent intelligence platforms, including application performance management, business intelligence, financial planning, and security operations, so storage signals participate in enterprise-wide, AI-driven operations rather than living in an infrastructure silo.
- Establish a continuous improvement loop where outcomes from autonomous decisions feed back into model training, policy refinement, and orchestration tuning, with measurable quality gains over time.

Who to involve:

- **Site reliability engineer:** Audits Everpure Fusion AI remediation module performance, validates real-time operations scoring integration for governance, and ensures autonomous action accountability is demonstrably enforced with a full audit trail
- **Financial analyst/CFO:** Defines cost-to-serve and unit economics financial metrics, ensures they are accurately tracked via Evergreen//One dashboards, and validates business unit alignment through financial reporting
- **Enterprise application architect:** Operates the Everpure Fusion compliance engine as the authoritative source for fleet-wide workload state and ensures continuous signals feed planning, governance, and executive reporting

Success indicators:

- Everpure Fusion compliance engine, Pure1 risk modeling, and intelligent orchestration tools operate as an integrated stack, with no manual data movement between layers.
- Pure1 and Everpure Fusion signals are consumed by at least two adjacent enterprise platforms (APM, BI, financial planning, SecOps, and so on).
- All workloads remain continuously in SLA compliance, with incident reports and immediate remediation when they don't.
- Continuous improvement metrics show measurable gains in autonomous resolution rate, forecast accuracy, and avoided incidents over consecutive quarters.

Everpure enablers:

- Everpure Fusion compliance engine for fleet-wide workload state and SLA drift alerting
- Pure1 adaptive risk modeling and automated forecasting
- Intelligent orchestration tools with cross-pillar, policy-enforced execution
- Everpure Fusion and Pure1 integration with adjacent enterprise platforms (APM, BI, financial planning, SecOps)

Maturity path: Level 5

Reaching Level 5 in application agility and automation means infrastructure has stopped being a bottleneck and started behaving like a business capability. By using this blueprint as a self-assessment, organizations can identify where manual processes and disconnected teams are slowing execution and build a clear path forward. Moving from the reactive, ticket-driven operations of Level 1 to the intent-based, autonomous orchestration of Level 5 takes deliberate work across technology, process, and governance, and it doesn't happen in a single initiative. But organizations that make the progression don't just deploy faster. They gain the ability to absorb business change without infrastructure delays, respond to market signals in near-real time, and compete on execution in a way that teams still managing provisioning queues simply can't.

Why Everpure

The platform built for tomorrow

Every recommendation in this blueprint is grounded in what the Everpure Platform actually delivers, not what sounds good in theory. That distinction matters because a lot of vendors describe a similar destination. The difference is whether the platform you're betting on was designed to take you there.

The EDC is a model for managing data as a unified, governed asset across on-premises, hybrid, and cloud environments. Achieving it requires a platform that operates at every layer: storage infrastructure, data intelligence, policy automation, and lifecycle management. Everpure is the only platform built to do all of that as an integrated whole.

What that means in practice

Continuous data intelligence

The semantics layer, powered by 1touch, continuously discovers, classifies, and contextualizes data across every environment, whether structured, unstructured, storage-as-a-service, mainframe, or cloud. It builds a live knowledge graph mapping the relationships between datasets, workloads, and policies. Every automated decision the EDC makes, whether placement, protection, or governance, is informed by the context required for accuracy. That's not a point-in-time scan. It's the difference between infrastructure that reacts and infrastructure that knows.

Policy-driven automation

Define intent once and the platform handles the rest. Provisioning, protection, life cycle, and compliance all run continuously against your policies, adapting as workloads change. Manual runbooks don't scale; policy does. Teams that make this shift typically cut manual operational tasks by 70–80%.¹

Evergreen//One: Always current, delivered as a service

The platform never falls behind, and there's no disruptive upgrades, migration projects, or planned downtime. 97% of systems sold in the last six years are still in service and current. When vendors talk about a modern platform, that number is what it actually looks like.

Guaranteed resilience

When something goes wrong, clean infrastructure is back in 24 hours, backed by a first-of-its-kind SLA guarantee. SafeMode Snapshots, always-on encryption, and continuous monitoring give you consistent resilience across the entire EDC, not just the systems someone remembered to protect.

The results organizations are seeing

These aren't theoretical outcomes. They're what happens when the platform is operating as designed and the maturity work in this blueprint has been completed. The numbers compound as you move up the maturity curve.

74%

reduction in TCO compared to legacy storage²

85%

average fleet utilization³

4–6 weeks

saved on project timelines

20x

faster AI processing⁴

The only platform that guarantees the outcome

Most infrastructure investments come with a lot of aspiration and very little accountability. Everpure is different in one specific way: it backs the outcome with SLAs. That level of commitment only makes sense if you’ve built the platform to deliver it consistently.

The EDC success blueprint gives you the vision. Everpure gives you the platform to make it a reality.

Get Started With Your Enterprise Data Cloud Success Blueprint

1 | <https://www.everpuredata.com/products/automation-orchestration.html>

2 | <https://www.everpuredata.com/company/newsroom/press-releases/pure-delivers-advancements-with-next-gen-flasharray-x-and-c.html>

3 | <https://www.everpuredata.com/products/automation-orchestration.html>

4 | <https://blog.everpuredata.com/purely-technical/20x-faster-inference-first-kv-cache-for-s3-and-nfs/>

Visit Our Website

800.379.PURE

