rackspace technology

Service Description

Rackspace DataOps



Table of contents

1 Overview 4
1.1 Key Features 4
1.1.1 Rackspace Elastic Engineering for Data 4
1.1.2 Data Pipeline Management – 24/7 Monitoring & Response (an add-on service) 4
1.2 Key Benefits 4
1.3 Supported Cloud Platforms 5
2 Rackspace Elastic Engineering for Data 5
2.1 Customer's Rackspace Elastic Engineering for Data: Experience 6
2.2 Rackspace Elastic Engineering for Data: Delivery Model 7
2.2.1 Rackspace Elastic Engineering for Data Pod Members 7
2.3 Onboarding 8
2.3.1 Customers transitioning from a Professional Services project 8
2.3.2 Customers with existing Customer created footprint 8
2.4 Continuous Development and Delivery 9
2.5 Hourly Consumption 10
2.5.1 Data Pipeline Management 10
2.6 Rackspace Elastic Engineering for Data Service Relationships (Roles and Responsibilities) 10
3 Data Pipeline Management 11
3.1 DPM Services 11
3.2 Onboarding 12
3.3 Service Level Agreement 12
3.4 Cloud Monitoring 13
3.5 Spheres of Management 13
3.6 Operations Runbook 14
3.7 Escalation to Rackspace Elastic Engineering for Data 15
3.8 Ticketing Process for Data Pipeline Management 16
4 Customer Success 17
4.1 Purpose 17
4.2 Objectives and Guiding Principles 18
4.3 Overall Partnership Governance Structure 18
4.4 Customer Interactions 19

4.4.1 Rackspace Elastic Engineering for Data 19
4.4.2 Rackspace Customer Portal 19
4.4.3 Operations Runbook 20
4.5 Pricing 20
4.6 Billing 21
5 Rackspace DataOps Service Delivery 21
5.1 Detailed Scope 21
5.2 The Data Pipeline Framework 23
5.3 Supported Platforms 23
5.3.1 Amazon Web Services 23
5.3.2 Google Cloud Platform 24
5.3.3 Microsoft Azure 24
5.3.4 Third party Technologies 24
5.4 Exclusions 25
5.5 Dependencies 25
Appendix 1: Support Request (Incident) Management and Resolution Process 26

1 Overview

Rackspace Technology delivers an organization-wide data management practice that improves the speed and accuracy of analytics through process automation and cultivating collaboration between data scientists, engineers, and technologists. This agile and automated, process-oriented methodology removes the operational burden of data platform management by:

- Automating tasks like data clean-up/prep, analysis and data integration
- Offloading engineering work including integrating with data sources, performing transformations, converting data formats, and writing or delivering data to its required destination
- Opening data access to more stakeholders within an organization utilizing industry best practices

With your data engineers no longer spending significant amounts of their time on these mundane tasks, they can shift focus toward higher value tasks like advanced machine learning, AI and real-time streaming. In turn, business operations and production can move faster when data teams are no longer strapped for resources.

1.1 Key Features

1.1.1 Rackspace Elastic Engineering for Data

Rackspace Elastic Engineering for Data provides access to a managed service that enables data lifecycle management to help increase operational efficiency and accelerate innovation. Rackspace Technology data experts work closely with Customer's team across one joint, agile team to collaboratively accelerate data initiatives, while passing on collective knowledge gained from multiple Rackspace Technology clients.

- **"Do with" Approach:** A Rackspace Elastic Engineering for Data Pod works in an agile, sprint-based model alongside your team.
- **Consistent Team:** No matter which tier of hours you use each month, you'll always work with the same pod that knows you, your environment and your business.
- **Flexible, Tiered Pricing:** Purchase "fractional access" to your pod via straightforward hours-based tiers and scale up and down should your business needs change.
- **Multi-faceted skill sets:** The Rackspace Elastic Engineering for Data Pod consists of an engagement manager, data architects and data engineers working together as one unit.
- **Ongoing Innovation:** Rackspace Technology leverages leading cloud platforms tools and frameworks to deliver progressive improvement, enablement and transformation.

1.1.2 Data Pipeline Management – 24/7 Monitoring & Response (an add-on service)

Our automated pipeline monitoring system is supported with reactive 24/7 response by our global Rackspace DataOps team. In addition to monitoring, alerting, and incident management, we provide troubleshooting, platform escalations, runbook-based resolution, and operational support. Data Pipeline Management is an optional add-on Service to Rackspace Elastic Engineering for Data.

1.2 Key Benefits

Rackspace Technology delivers process maturity, agile delivery and world-class support so your team can make business and IT decisions that will drive the most value to your customers and have the greatest impact to your bottom line. Rackspace DataOps delivers results that accelerate innovation.

- A managed end-to-end analytics pipeline that reduces operational burden and allows you to focus on business-critical activities
- Automation and accelerated adoption of platform capabilities
- Expert direction based on the latest innovations in data management
- Operational data management principles that deliver a repeatable approach to data lifecycle management
- Collaboration around data lifecycle and governance that establish shared responsibility for successful delivery

1.3 Supported Cloud Platforms

The services described in this document apply to the following public cloud platforms:

- Amazon Web Services
- Google Cloud Platform
- Microsoft Azure

2 Rackspace Elastic Engineering for Data

Rackspace Elastic Engineering for Data is a managed support model that provides on-demand access to data experts via ongoing sprints to continually remediate, innovate, evolve, and improve customer's environment(s). Rackspace Elastic Engineering for Data is delivered as tiers of service that include a number of hours per month.

This flexible, collaborative approach to services allows customers to dynamically change scope and priorities based on business needs while collaborating side-by-side with highly skilled Rackspace Technology data experts to architect, build, optimize, and maintain the Customer's cloud analytics environments.

The table below provides a sample of the types of services provided with Rackspace Elastic Engineering for Data. Tasks, scheduling, and priorities are coordinated between Rackspace Technology and the Customer.

Service	Description
Data Strategy	• Work with customer to prioritize new use cases according to business needs and solution complexity, define amendments to reference architecture, and create an implementation roadmap
Data Platform Architecture	• Continuous data platform architecture and engineering evolution of the customer environment focusing on improving performance and cost through continuous optimization and new features
Data Lifecycle Automation	• Expand and renew existing automation throughout the data lifecycle across production and sandbox environments, from ingestion of CI/CD pipelines for data pipeline and big data jobs, to schema updates and tiering of data
Operational Excellence	• Build and maintain monitoring, auditing, logging and tracking across all assets in the data platform architecture and provide continuous

Service	Description
	improvement to operational posture by helping reduce overall error rates and issues caused by operating data platforms at scale
Data Quality	• Improve, validate and integrate tests into existing and new data transformations to help clean, prepare and maintain the quality of the data
Data Security	• Build and enhance security controls from ingestion to processing and storage at the data platform layer, with consideration of Role Based Access Controls, access permissions, data tokens, and building a governance framework to audit and review access to data sets
Observability	• Provide on-going advice and recommendations based on profiling current workloads and activities
Schema & Performance Management	• Manage metadata and data catalog with change controls and performance management of data workloads and queries
Runbooks	• Create and update Customer provided, and Pod created runbooks to support managing, troubleshooting and remediating a Customer environment

2.1 Customer's Rackspace Elastic Engineering for Data: Experience

With Rackspace Elastic Engineering for Data, customers purchase monthly hours that can be used to engage with a team of Rackspace Technology cloud data experts assigned to a customer's account. Rackspace Elastic Engineering for Data is organized into a group known as a pod that operates according to the following principles:

- Pods offer a shared resource model
- Sprint-based project methodology with emphasis on partnering with the Customer's internal analytics, development and operations teams
- Emphasis on enablement, automation, and transformation
- Build a practice of ongoing innovation, rather than focus on break-fix
- Execute small projects that deliver continuous, incremental improvement
- Help Customers evolve into a data driven business

Customers are not locked into long-term contracts and are given the opportunity to adjust their tier up or down to match needs and budget constraints.

As opposed to one-time pro-serve engagements and traditional support, Rackspace Elastic Engineering for Data has the benefit of continuity and familiarity. The pods deeply understand the customer's business requirements and data environment, which makes them more efficient and aligned with the expected outcomes.

2.2 Rackspace Elastic Engineering for Data: Delivery Model

A Rackspace Elastic Engineering for Data Pod possesses the following characteristics:

- **Diversified group of experts:** A pod includes data experts, architects, engineers, and an engagement manager who work together as a unit.
- **Consistent team structure:** Pod members have been working together long-term across multiple clients and projects.
- **Shared knowledge and resources:** Pod members create and share resources, platform tools, and frameworks with other pods. Customers benefit from this collective knowledge.

2.2.1 Rackspace Elastic Engineering for Data Pod Members

Your Rackspace Elastic Engineering for Data Pod may consist of the following members:

Engagement Manager (EM): The EM is the Customer's primary point of contact and is responsible for the success of the customer engagement. The EM oversees tier-hours consumption, customer requests, and acts as a point of escalation. The EM also acts as a Scrum Master, assisting the team in:

- Project planning
- Resource coordination
- Establishing project timeframes and quarterly goals
- Scheduling and leading daily standup meetings, sprint reviews, monthly retrospectives, and quarterly business reviews (QBR)
- Providing work accomplished, budget, and planning delivery reports

Lead Data Architect (DA): The DA owns the successful delivery of a solution that meets the Customer's technical objectives. The DA provides program level technical guidance and advice, co-develops data strategy and roadmap and best practices consultation and owns the overall technical management and performance of the Rackspace DataOps Framework for the Customer.

Senior Data Engineer (SDE): The SDE contributes technical expertise in one or more domains in the context of the larger solution. The SDE collaborates with the Data Architect and peers to deliver a complete data solution.

Data Engineers (DE): The DEs assist the Customer with development, process and pipeline improvements working closely with the Data Architect and interfacing with the operations engineers on escalations and high priority activities.

Business Intelligence (BI) Developer: The BI developer focuses on improvements, updates and building new dashboards and reports to support the growing and changing business and operational reporting needs.

Data Operations Engineers (DOE): For Customers consuming Data Pipeline Management, their data platform will be supported around-the-clock by our global Rackspace DataOps team. The DOEs will monitor customer environments, respond to alerts, and resolve issues based on defined runbooks.

Rackspace Technology shall always ensure there is an appropriate number of personnel to perform the Services during the supported hours.

Rackspace Technology shall ensure that all its personnel:

- Are appropriately experienced, qualified, and trained to perform the Services.
- Perform the Services with reasonable skill, care, and diligence.
- Cooperate with Customer personnel to the extent reasonably necessary to provide the Services.
- Are authorized to perform Customer-agreed activities on the supported Services.

2.3 Onboarding

At the beginning of the Rackspace Elastic Engineering for Data engagement, Rackspace Technology schedules a project kickoff call with the Customer. For new customers, a walk-through of the Customer Portal is offered to help Customers become familiar with Rackspace Technology ticketing services and notifications. All onboarding work is delivered as Rackspace Elastic Engineering for Data time and uses hours from the Customer's selected Rackspace Elastic Engineering for Data tier.

The onboarding team can consist of:

- Customer Success Manager
- Engagement Manager
- Lead Data Architect
- Data Engineers

2.3.1 Customers transitioning from a Professional Services project

- **Deliverable Transfer:** The Professional Services team involves the Rackspace DataOps SME during the design phase and creates the necessary architecture diagrams, design documents, automation and runbooks that will be leveraged for steady state operations of the environments.
- **Kick-off:** The on-boarding team provides a kick-off call to officially transfer all deliverables from the Professional Services team to the Rackspace Elastic Engineering for Data Pod, who will then take over the responsibility of updating the design documents, architecture diagrams, automation and code.
- **Discovery:** The Customer, the Professional Services team, and the Rackspace Elastic Engineering for Data Pod will take part in discovery sessions and technical deep dives to transfer customer specific processes, custom workflows and code and data management overviews.
- **Transition:** Responsibility of the supported environments is transitioned to the Rackspace Elastic Engineering for Data Pod and the engagement begins in earnest.

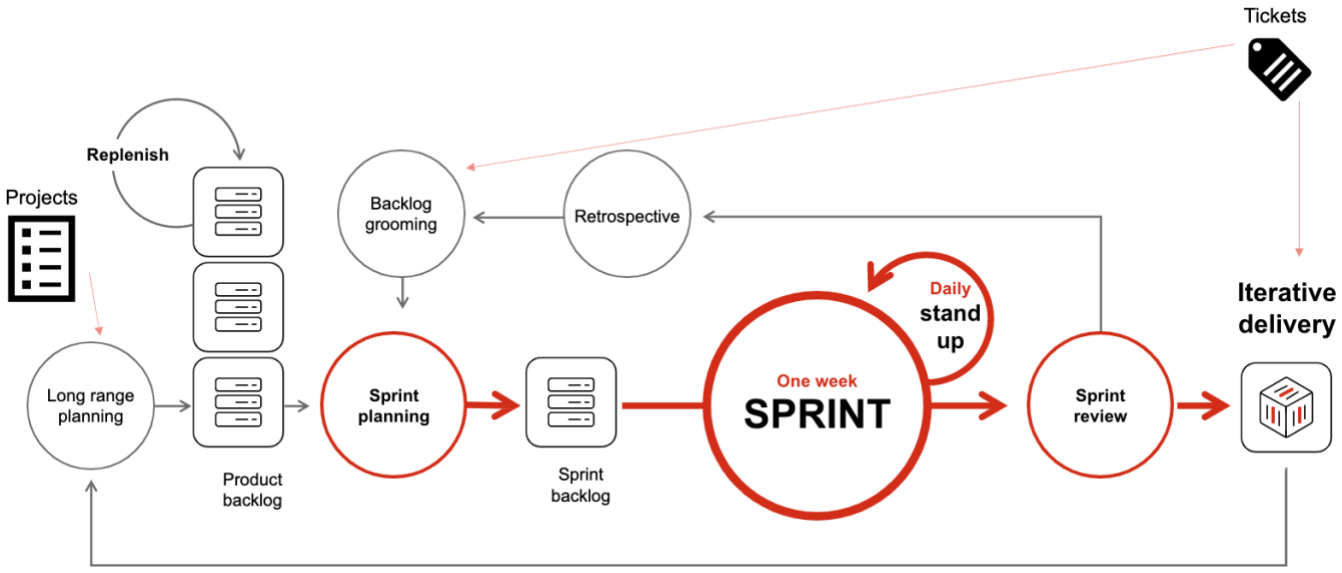
2.3.2 Customers with existing Customer created footprint

- **Kick-off**: The on-boarding team provides a kick-off call to begin the official process of transferring the responsibility of the management the environment from the customer
- **Discovery**: Several discovery calls will be scheduled with the Customer to provide detailed accounting and documentation of the existing environments, process and policies in place. Customer provides all existing documentation, architecture diagrams, workflows, automation and code
- **Assumption**: Once the onboarding team receives all the supporting documentation and code for the environments and gains the required access to all the materials and environment, the team performs a gap assessment of runbooks, automation and documentation required to effectively support the environment
    - The minimum required documentation and runbooks to start the transition will be created by the on-boarding team
    - The remaining documentation, runbooks, automation and code required will go into the backlog and continuously improved and incorporated over the proceeding months of service

- **Transition:** Responsibility of the supported environments is transitioned to the Rackspace Elastic Engineering for Data Pod and the engagement begins in earnest

2.4 Continuous Development and Delivery

The following diagram illustrates how Rackspace Elastic Engineering for Data and Customers partner in the continuous development and delivery model.



Engagement Delivery Framework: The Rackspace Elastic Engineering for Data delivery framework provides an agile process which closely follows the principles of scrum. This approach provides ongoing visibility to key stakeholders, which allows required course corrections to be made throughout the project. Rackspace Technology begins the engagement with a long-range planning session, where the project is divided into achievable work items that can be delivered and demonstrated in weekly iterations.

Backlog and Sprint Planning: Before each iteration (known as a *sprint*), Rackspace Technology works with Customer to document and review a backlog of user stories. This planning meeting is known as backlog grooming. At the start of each sprint, Rackspace Technology and the Customer participate in another sprint planning meeting where new stories for the sprint are selected from the top of the backlog. As much as possible, any unknown details regarding the selected stories will be resolved during this meeting. If a story requires additional definition, analysis, or design work, then an alternate task to accomplish that work during the sprint can be selected in place of the ambiguous story.

During the Sprint: A sprint begins on the day of the sprint planning meeting and ends the day before the next planning meeting. During the sprint, Rackspace Technology data engineers conduct all analysis, design, engineering, and testing activities. The team meets each day for a standup meeting, where each member discusses their previous day's accomplishments, their plan for the day and any issues that are blocking progress.

Sprint Demo and Retrospective: The end of each sprint is celebrated with a demonstration to key stakeholders of the functionality completed during the sprint. The end-of-sprint demonstration allows the Customer to

continually monitor the progress of the project. The team also holds a retrospective meeting, where lessons from the completed sprint are used to improve future sprints.

Customer Acceptance: Rackspace Technology is responsible for the quality and performance of its resources. Rackspace Technology adheres to its standard Engagement Delivery Framework and during the sprint planning meeting, stories and acceptance criteria are reviewed with the Customer. At the end of each sprint Rackspace Technology reviews the completed stories with the Customer's stakeholders. Stories are marked complete if they meet the agreed upon acceptance criteria. Stories are marked incomplete and worked during the next sprint if they do not meet the agreed upon criteria.

If a story does meet the agreed upon acceptance criteria, and the Customer wants to make changes to that story, an additional story is created, added to the backlog, and prioritized by the Customer and incorporated into a later sprint.

2.5 Hourly Consumption

Customers will be invoiced each month, based on the tier purchased.

Unused hours do not carry over to the next month, and no credits or refunds are available for unused hours. On the service start date, the Customer starts at the purchased tier. Customers can change tiers subject to the following restrictions:

- Tier change notices must be communicated via a ticket submission to Rackspace Technology.
- Tier changes require 30 days advance written notice.
- Customers cannot give notice to reduce to a lower tier within 30 days of increasing to a higher tier.
- Any notice received during a 30-day period will be effective on the first day of the month after an additional month. For example, if a customer notifies Rackspace Technology of a tier change on March 15, the notice is effective on May 1.

2.5.1 Data Pipeline Management

Customers who add Data Pipeline Management to Rackspace Elastic Engineering for Data receive 24x7x365 monitoring and alert response. For more details on the offer and response time SLA' s, refer to Data Pipeline Management.

2.6 Rackspace Elastic Engineering for Data Service Relationships (Roles and Responsibilities)

The following table identifies who is responsible and accountable for specific engagement activities:

- **Responsible** means, for the purposes of the RACI matrices, the party who carries out the process or task assignment and is responsible to get the job done.
- **Accountable** means, for the purposes of the RACI matrices, the person who is ultimately accountable for the process or task being completed appropriately.
- **Informed** means, for the purposes of the RACI matrices, those who receive output from the process or task, or who have a need to stay informed.

Activity	Customer	Rackspace Technology
Provide Rackspace Technology with a list of the Customer's engineers that will use its services.	Responsible	Accountable
Periodically update list of the customer's engineers that will use Rackspace services.	Responsible	Accountable
Train the Customer's users, clients, and suppliers to follow established Rackspace procedures to report incidents or make requests.	Informed	Responsible
Provide Rackspace Technology engineers with full administrative access to all Cloud accounts supported by the pod.	Responsible	Accountable
Make staff available to work with Rackspace resources during the handling of an incident or request.	Responsible	Informed
Provide Rackspace Technology all documentation, information, and knowledge as requested in relation to the Cloud environments supported by the pod.	Responsible	Accountable
Provide customers timely reports related to off hours incidents and post-mortem of any infrastructure related incidents that occur in accounts managed by the assigned pod.	Informed	Responsible

See Appendix 1 for the Support Request (Incident) Management and Resolution Process.

3 Data Pipeline Management

Built for Customer's critical workloads, an always vigilant 24x7x365 team of data operations engineers monitor and respond to alerts in Customer's environment. Data Pipeline Management (DPM) is an opinionated deployment of cloud services monitors that act as an extension of Customer's operational team by providing automation where desired and human actions for activities that are beyond automation.

With Rackspace Technology acting as an extension of your 24x7 operational team, you can focus on what matters most to your business.

3.1 DPM Services

DPM provides real-time monitoring and alerting for your data analytics environment. An expert team of Rackspace DataOps engineers monitor and respond to any alerts in your environment. Rackspace Technology has the breadth of knowledge required to operate the most complex data pipelines on a global scale. Combined with the Rackspace Elastic Engineering for Data service, you'll gain the peace of mind that comes with having an environment that's cared for by a team of Rackspace DataOps engineers that help drive continuous innovation.

DPM includes:

- 24x7x365 Monitoring and Alert resolution including triage and root cause analysis of data pipeline failures to maintain healthy movement of data within the platform based on Runbooks
- Support for PaaS and Serverless data services focused on resolving and remediating issues and performing activities as outlined in the Runbooks
- Opinionated deployment of cloud services monitoring using known best-practices
- Automation to improve resolution timelines
- Guidance for a default list of alerts and thresholds, which can be customized to meet customer needs
- On-going support to adjust monitoring as customer environment grows

Note: DPM requires Rackspace Elastic Engineering for Data, and Optimizer+ or Optimizer+ with Platform Support as prerequisites.

3.2 Onboarding

Onboarding activities are critical to ensuring that the right level of monitoring is applied to workloads based on their criticality. The DPM onboarding activities are either carried out as part of the Rackspace Elastic Engineering for Data onboarding if purchased together or separately if purchased later. At the beginning of the engagement the assigned Rackspace Elastic Engineering for Data Pod works with the customer to complete the following onboarding activities:

1. Meet with the Customer to advise on a monitoring strategy that meets the customer's requirements.
2. Create and update Customer provided and Pod created runbooks to document cloud environment monitors, thresholds, and responses.
3. After Rackspace Technology and the Customer agree to the runbook, Rackspace Technology deploys the monitors into the Customer's cloud environment.
4. Rackspace Technology works with the Customer to compile a list of additional runbooks that are applicable for the customer environments. Rackspace Technology will take over and update any existing runbooks, or created new runbooks as needed.

3.3 Service Level Agreement

For Customers with DPM, Rackspace Technology commits to the following initial response times for alerts.

Severity	Definition	Initial Response Time
Emergency	Business-critical system outage / extreme business impact. Ex: Operational pipelines feeding real-time systems with immediate business impact	15 minutes
Urgent	Business-critical system impairment / with significant business impact Ex: A key component of the Analytics stack, the pipeline, warehouse or BI service impact for end users of the applications or data	60 minutes

Severity	Definition	Initial Response Time
High	Production system impairment with moderate business impact. Ex: Jobs scheduled 6 hours apart and higher that fail and need to be investigated and remediated	4 hours
Normal	Minimal business impact. Ex: User permissions requests, alerts with soft thresholds	12 hours
Low	General monitoring information, questions, and guidance	24 hours

3.4 Cloud Monitoring

Data Pipeline Management leverages cloud native monitoring systems to deliver the monitoring and management of the data pipelines and data platform services. This is CloudWatch on AWS, Azure Monitor on Azure, and Cloud Monitoring in GCP.

The cloud monitoring service available with AWS, Azure, or GCP is the primary monitoring system used by Rackspace DataOps support teams. While the cloud monitoring service will be available to Customer accounts at all service levels, Customers using the service can opt to have Rackspace Technology respond to generated alarms. Rackspace Watchman provides a combination of the systems responsible for creating Rackspace Technology support tickets from the alarms and the certified experts who take the actions necessary to mitigate the indicated alarm conditions 24x7x365.

Cloud monitoring alarms are configured to trigger when the value of a specified metric deviates from its expected value for a specified period of time. For example, if a job fails or reaches a threshold on duration of execution, the alarm is configured to send an alert to Rackspace Watchman to generate a ticket for further investigation by Rackspace Technology.

3.5 Spheres of Management

The following list describes the core set of activities that Rackspace Technology performs as part of the Data Pipeline Management service component of Rackspace DataOps.

- Data Pipeline monitoring and management
- Configuration and setup of services as defined in runbooks for repeatable processes. Ex: setup of sandbox environments, test data preparation
- Basic user management (including create, delete, and password reset)
- Performance troubleshooting and issue resolution within the environment
- Platform related escalations
- Data Health monitoring

- Response and remediation of alerts
- Resource and security management (roles and permissions)
- Working with stakeholders for data related issue resolution
- Configure monitoring based on well-established best practices and principles using Cloud vendor tooling, including:
    - GCP Cloud Monitoring
    - AWS CloudWatch
    - Azure Monitor – Metrics and Azure Monitor - Log Analytics & Application Insights - Web Availability Test.
- Follow alert remediation procedures as outlined in the Runbook.
    - The Rackspace Technology response is bound by the SLAs documented in the Service Level Agreement section.
    - Rackspace Technology restores the service to an available or degraded-but-functional state either by following the runbook, taking appropriate actions, or by escalating to the customer.

3.6 Operations Runbook

Rackspace Technology shall work with Customer to create and update operations runbook. The runbook defines the Data Pipeline Management standard operating procedures for the response to monitoring alerts, management and response to pipeline failures and data ingestion failures in the environment and includes custom escalation procedures in accordance with best practices and Customer's business needs. These customer runbooks are designed to present the right information, at the right time, to Data Pipeline Management. It's important to be able to respond quickly and effectively to service disruptions. Providing relevant and focused guidelines to Data Pipeline Management increases the availability of customer solutions.

Rackspace Technology shall:

- Develop an operations runbook in accordance with the transition plan:
    - The operations runbook describes the procedures to be used to perform the Services.
    - The operations runbook conforms to Rackspace Technology policies and, where reasonably possible, Customer's policies.
    - Rackspace Technology and Customer shall jointly use the operations runbook to enable close cooperation and communication between the parties.
    - The operations runbook includes checkpoint reviews, testing, acceptance, and other procedures for Customer to assure the quality of Rackspace Technology performance.

- Perform the Services in accordance with Customer-approved operations runbook.

The operations runbook contains the following items, at a minimum:

- AWS Accounts, Azure Subscriptions, or GCP Projects in scope
- Customer Contact Call Order
- Rackspace Contact Call Order
- List of Monitors
- List of Response Procedures
- Pre-approved actions
- Supporting Documentation

- Cloud Dashboards

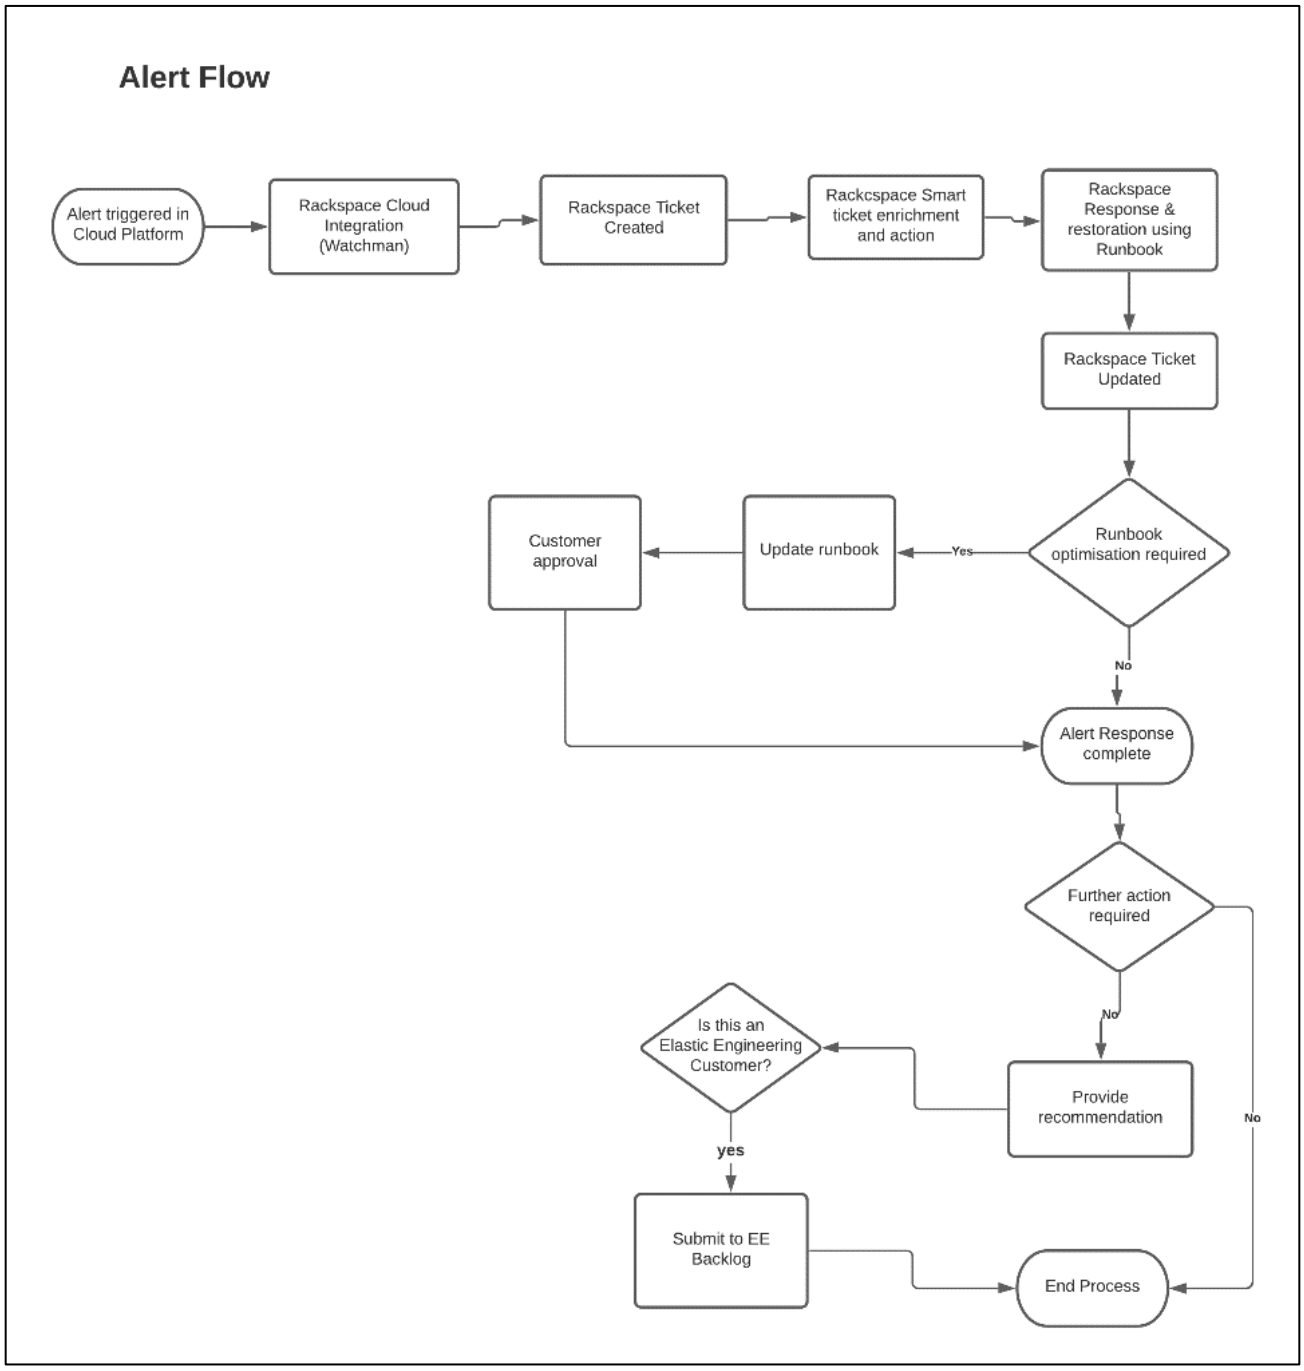
3.7 Escalation to Rackspace Elastic Engineering for Data

When an alert response indicates that an escalation is required, Data Pipeline Management will escalate to the respective customer Pod. The following list include some of the activities that typically require close collaboration or escalation to Rackspace Elastic Engineering for Data.

- Onboarding
- Monitoring changes
- Updated and new services, jobs and data sources
- Long-term design changes and recommendations
- Emergency escalation for system-down events.

3.8 Ticketing Process for Data Pipeline Management

The following diagram illustrates the stages of alert handling where Rackspace Technology is involved.



Stage 1: Monitoring is performed by the Cloud monitoring tooling. For monitors that require action by Rackspace Technology, Rackspace Technology creates alert policies that are ingested by Rackspace Watchman.

Stage 2: Rackspace Watchman generates tickets using the determined priority level and then passes the ticket to the Rackspace Smart ticket system.

Stage 3: The Smart ticket system enriches the alert by collecting diagnostic information about the alert. The diagnostics can be as simple as collecting audit logs close to the alert time or running diagnostics on the underlying operating system.

Stage 4: At the same time the Smart ticket system collects diagnostics information about the ticket, the ticket is visible to the Rackspace Technology support team who can start taking actions documented in the Runbook.

One of the primary ways that Customer can interact with Rackspace Technology is by creating a ticket in the Rackspace Customer Portal (https://login.rackspace.com/login). After login, click the **Tickets** button from the menu to create a new ticket or view an existing ticket. Rackspace Technology automated systems also create tickets for events on Customer's account that require either Customer's attention or the attention of a Rackspace Technology employee.

Detailed ticketing process will be outlined during the on-boarding process.

Rackspace Technology shall:

- Respond to tickets created by Customer in the Rackspace Customer Portal.
- Respond to tickets created by the Rackspace Technology automated system for events on Customer's account or accounts that require the attention of Rackspace Technology or Customer.

Customer shall:

- Primarily interact with Rackspace Technology data operations team by creating a ticket in the Rackspace Customer Portal.
- Respond to tickets created by the Rackspace Technology automated system for events on Customer's account that require attention.

4 Customer Success

Depending on Customer's service level, this section covers the following areas of Customer Success:

- Purpose
- Objectives and Guiding Principles
- Overall Partnership Governance Structure

4.1 Purpose

The purpose of the governance structure is to ensure that guiding principles, objectives, structures, operating guidelines, methods, and measures for implementing effective governance are clearly defined and consistently implemented.

The governance function is subject to change based upon Customer's decisions on target organization design. The Rackspace DataOps service was designed based on feedback from Rackspace Technology Customers with a flexible engagement model to not only provide support coverage but also keeping their data platform up to date and working closely with customer teams as data is a critical component to any organization.

4.2 Objectives and Guiding Principles

This section describes the structure of the relationship model, with principles of partnership that define the levels in which Rackspace Technology and Customer interact and align resources. It also defines a framework for innovation and growth that enables the relationship to evolve as Customer's business changes, so that the partnership constantly works toward value improvement.

The governance model is designed to achieve the following guiding principles:

- A multi-level service structure to provide end-to-end need for managing and innovating a data platform.
- Provide a shared resource model and a build-with model to promote trust through transparency and bilateral communication.
- Maintain the strength of the relationship between Customer and Rackspace Technology as a critical success factor for both.
- Emphasize the orchestration of people, processes and technology to deliver trusted, high-quality data to whomever needs it by establishing a shared organization and structure to streamline day-to-day management and administration.
- Develop a process framework for the entire data lifecycle that incorporates DevOps style deployments, agile development and statistical process controls to apply towards people, process and technology.
- The goal of Rackspace DataOps is to create predictable delivery and change management of data, data models and related artifacts.
- Ensure overall monitoring of contract performance on service levels, financials, deliverables, and customer satisfaction.
- Ensure that potential issues are investigated, resolved and/or, if necessary, escalated.

4.3 Overall Partnership Governance Structure

The purpose of governance is to establish effective means for managing the delivery of Services and innovation as determined by Customer's objectives and this Service Description. This is a joint framework managed equally by Rackspace Technology and Customer.

Regular Touchpoints

- Weekly touchpoint calls for review and planning of activities and key events for the week
- Create a roadmap and backlog of tasks for any co-development efforts
- Monthly review and planning meetings covering work accomplished, ticket flow, next month planned activities
- Quarterly Business Reviews (QBR)

Reporting

- Month end report covering:
    - Allocated and completed work items
    - Backlog and work items for next month
    - Time tracking and ticket statistics

- Quarterly reports:
    - Highlight key achievements and issues

    - Recommendations on tiering and development effort changes
    - Consolidated metrics from monthly reports for the quarter

- Data Pipeline Management reports:
    - Ticket statistics that include the following
        - Number of tickets opened, resolved, escalated
        - Ticket SLAs with response times
        - Ticket resolution times
    - Monitoring recommendations and pipeline improvements
    - Metrics and job statistics

4.4 Customer Interactions

This section describes the customer interactions for the Service, which includes the following:

- Rackspace Elastic Engineering for Data
- Rackspace Customer Portal
- Operations runbook
- Ticketing process

4.4.1 Rackspace Elastic Engineering for Data

Customer's Rackspace Elastic Engineering for Data Pod will be available via direct access thru email and either a customer communications tool or thru Microsoft Teams on a shared customer channel. Tasks, and the scheduling and prioritization of those tasks will be defined by Customer and Rackspace Technology through creation of a backlog in Jira or such other project management system as Rackspace Technology may later designate.

4.4.2 Rackspace Customer Portal

Rackspace Technology provides Customer with access to the Rackspace Customer Portal (https://login.rackspace.com/), a consolidated login portal.

The Rackspace Customer Portal is the primary customer interaction point for Rackspace Technology Customers. It provides Customers with capabilities based on their service level, including the following:

- Create tickets
- Accept contracts
- View and pay bills

Customer can also manage fine-grained permissions for their account users.

The Rackspace Customer Portal focuses on Customer's ease of use by:

- Improving Customer account experience.
- Adding self-service features that reduce support volume.

Rackspace Technology shall:

- Provide Customer with access to the Rackspace Customer Portal (https://login.rackspace.com/).

Additional portal and service access as required will be provided by the on-boarding team.

4.4.3 Operations Runbook

Rackspace Technology and Customer can develop standard operating procedures (SOPs) to ensure that Rackspace Technology and Customer are aligned on common tasks. Rackspace Technology and Customer can define SOPs during the onboarding process or at any time during the term of service. Rackspace Technology can also leverage existing Customer defined SOPs that can be transitioned for Rackspace Technology to manage and update during the on-boarding process.

All change management activities required by the customer are to be documented and shared with Rackspace Technology to be incorporated as part of the SOPs.

Examples of the most common SOPs include:

SOP	DESCRIPTION
Monitoring events and remediation	Documentation of comprehensive details of all monitoring events that have been implemented to the Customer's environment and remediation steps for those events.
Escalation procedure	Documentation of escalation procedures for both the Customer team and the Rackspace DataOps team based upon various types of events.
Pipeline Failure	Process of responding to pipeline failures that could require restarting, after data validation and remediation of processing step errors. Resolutions could also include escalating to the Customer engineering team or the Rackspace DataOps development support team.
Sandbox environments	Set of procedures for creating or destroying sandbox environments as required for development and research teams

4.5 Pricing

Pricing consists of 2 layers with multiple tiers to choose from per the workload management requirements and Customer's on-going development needs.

- Rackspace Elastic Engineering for Data
    - A utilization-based model with a fixed set of hours per month
- Data Pipeline Management
    - Consumption is evaluated every quarter to place Customers in the right Tier for their workloads
    - Tier adjustments are discussed during QBRs

4.6 Billing

Rackspace DataOps is billed as a subscription-based service based on the Tier pricing and services subscribed which could be Rackspace Elastic Engineering for Data alone or with Data Pipeline Management added. Customer will have a line item on their monthly Rackspace Technology bill that lists the Rackspace DataOps subscription.

For Customers where it applies there is a one-time setup fee that pays Rackspace Technology for onboarding the Customer's environment and ensuring that it is ready for the Rackspace Elastic Engineering for Data Pod and the Global Rackspace DataOps Team to manage. The setup fee is included on Customer's first monthly Rackspace Technology bill.

5 Rackspace DataOps Service Delivery

This section describes the spheres of management and exclusions for Rackspace DataOps.

5.1 Detailed Scope

The list below provides an example of activities performed.

Monitoring and Alerting

- All services and activities will be monitored through platform native monitoring services.
- Thresholds defined based on mutually agreed items to create alerts across services for various metrics.

Logging and Auditing

- All monitoring metrics, service logs, audit logs and security logs pushed to platform logging services for deeper levels of alerting based on additional analytics.
- Provide quick access to logs for troubleshooting and performing analytics on applications and service usage and identify patterns.
- Auditing for service access, security and compliance.

Performance

- Enable Auto scaling and auto-shutoff features where available for services.
- Workload management and data distribution to support performant queries and prioritize workloads without cost overruns. Reserved vs on-demand utilization of resources.
- Profile and continuously improve transformation jobs and ingestion job performance.

Infrastructure as Code

- Implement and improve automation to reduce errors and simplify processes to create repeatable and reusable patterns.
- Sandbox environments for developers with ability to quickly spin up and spin down environments with testing data during the development process.

CI/CD Pipelines

- Manage, support and advice around release management, change management and CI/CD process for data pipelines and database schemas.
- Integrate all job code and schema with centralized repositories.
- Schema managed via version control and schema updates published thru pipelines.
- Create and improve CI/CD pipelines for analytics jobs with automated and gated deployments providing for code progression thru environments.
- Automation to improve efficiency and reduce costs in repeatable and resource heavy tasks.

Analytics Pipelines

- Monitor and manage and troubleshoot all analytics pipelines, Datawarehouse, data model deployments and related data services.
- Improve the overall efficiencies of the jobs, simplify deployments and automate Database Schemas, Stored Procedures, Triggers etc.
- Pipeline Artifacts such as Integration Runtimes, Linked Services, Activities, Data Flows, Pipelines.
- Hadoop/Spark artifacts such as spark jobs, hive/pig scripts, Jupyter Notebooks.
- Semantic Models.

Security

- Encryption at all layers, Transparent Data Encryption for storage layers and databases, always encrypted connections for data in transit.
- Fully defined access polices per service and environment getting down to the granular levels to support multiple domains, granular permissions with row level access controls for data marts.

Lifecycle and Governance

- Agile development methods and integrated with customer sprints.
- Automate the movement of test data including masking and tokenizing data where required.
- Support quick loads and re-loads of test data.
- Be the focal point in the organization governing the use of the systems to managing the data intake and updates to existing pipeline process and new data flows.
- Provision and audit access control policies to manage the use and consumption of data.

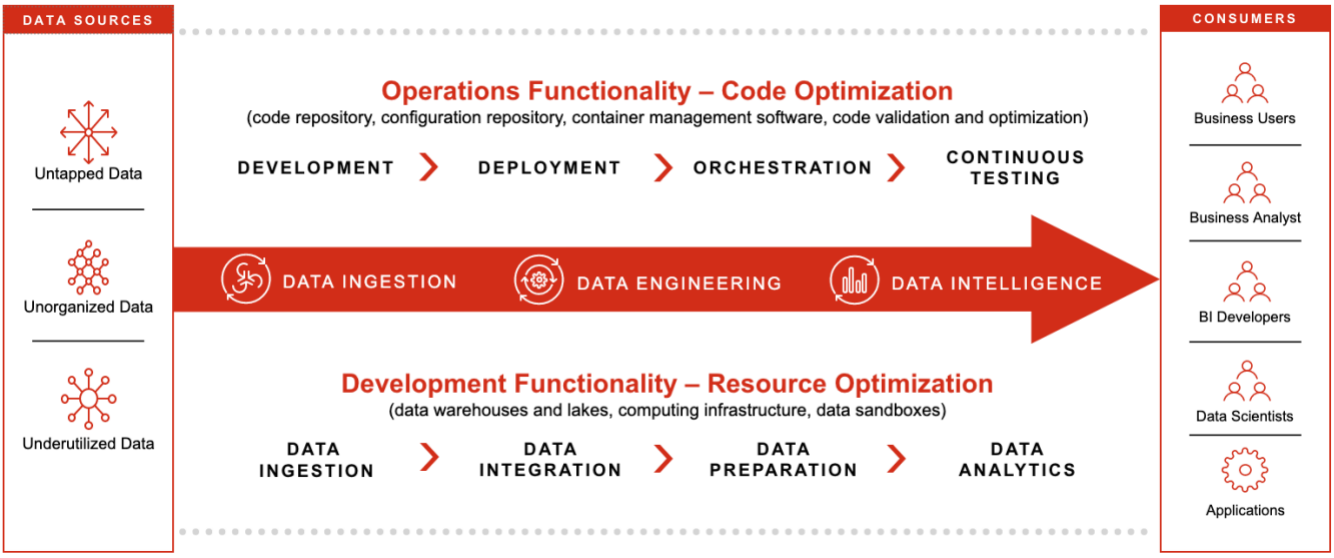
Collaboration

- Create a collaborative work environment between Rackspace Technology and Customer resources.
- Share and communicate best practices and agile and lean methodologies.
- Create repeatable process and frameworks to build the next set of use cases.
- Inculcate Rackspace DataOps process and methodologies to improve collaboration and communication across various teams within the customer organization.

5.2 The Data Pipeline Framework

This section provides a guide to applying Rackspace DataOps across code and resource optimization.

The data pipeline path runs from source to consumers and includes ingestion, engineering, and analysis in order to transform the data into useful intelligence. Data Development Functionality consists of the resources used to build the data pipeline—including replication, ETL, data catalogs, data quality, lineage tracking, data science, and business intelligence (BI) tools. Data Operations Functionality represents the components used to manage and optimize the creation and deployment of data pipelines.



5.3 Supported Platforms

This section provides information about the various third-party Cloud platform services in the Big Data and Analytics space and other third-party technologies supported by Rackspace DataOps specifically for the Customer.

5.3.1 Amazon Web Services

The following is a list of core services that will be supported on AWS.

- Amazon Athena
- Amazon CloudSearch
- Amazon Elasticsearch Service
- Amazon EMR
- Amazon Kinesis
- Amazon Managed Streaming for Apache Kafka
- Amazon Redshift
- Amazon QuickSight
- AWS Data Exchange

- AWS Data Pipeline
- AWS Glue
- AWS Lake Formation

5.3.2 Google Cloud Platform

The following is a list of core services that will be supported on GCP.

- BigQuery
- Dataproc
- Dataflow
- Looker
- Pub/Sub
- Cloud Data Fusion
- Cloud Composer
- Data Catalog
- Dataprep
- Google Data Studio
- Cloud Bigtable

5.3.3 Microsoft Azure

The following is a list of core services that will be supported on Azure.

- Azure Synapse Analytics
- Azure Analysis Services
- Azure Data Explorer
- Azure Data Lake Storage
- Azure Data Share
- Azure Databricks
- Azure Stream Analytics
- Data Catalog
- Azure Data Factory
- Data Lake Analytics
- Event Hubs
- HDInsight
- PowerBI
- Azure Purview

5.3.4 Third party Technologies

The following is a list of other third-party, supported technologies.

- Snowflake
- Databricks
- Tableau
- Hadoop ecosystem
- Spark

- Kafka

5.4 Exclusions

The following exclusions apply to Rackspace DataOps service support:

- **Professional Services SOW:** Any new and significant data engineering and development tasks arising from support activities or requested separately by the customer are treated as one-time events that require an additional Professional Services SOW. Customers can work with the Professional Services team to determine the scope and additional cost which would be treated separately from this service block.
- **Large scale Data Migration:** Complex and large-scale data migration services are not included in this offering. Any migration needs/requirements that the customer has must be addressed in a Professional Services SOW.

5.5 Dependencies

If Customer's platform environment is on Rackspace Managed Public Cloud, they must at a minimum have Rackspace Optimizer+ service block in order for Rackspace DataOps teams to have the appropriate access and monitoring solutions in place.

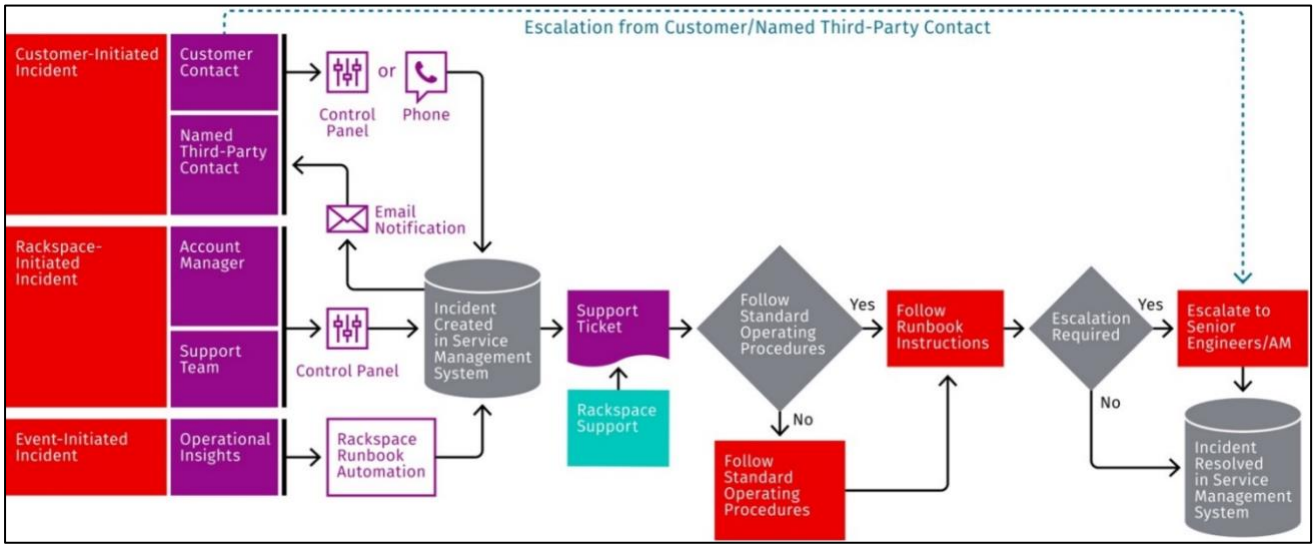
Access: Customer must provide Rackspace Technology with the appropriate access into their environment for the necessary monitoring and access points to be created to provide the services described. During the onboarding process, Rackspace Technology will work with Customer on getting both integrations in place.

On-Boarding: The on-boarding process could take anywhere from 1 month to 3 months depending on the complexity of the customer environment and the number of runbooks that will need to be created.

Knowledge Transfer: Detailed knowledge transfer is required in situations where a customer has a much more mature data platform with complex processes and multiple environments.

Appendix 1: Support Request (Incident) Management and Resolution Process

The following diagram shows the Support Request (Incident) Management and Resolution Process:



This Service Description is not a contract, and nothing in this Service Description is or may be construed as an offer or legal obligation on the part of Rackspace Technology or its Representatives. This Service Description may change from time to time, in the sole discretion of Rackspace Technology. Capitalized terms used but not otherwise defined in this Service Description will have the meaning given to them in the Rackspace Technology Master Services Agreement (presently found at the following URL:https://www.rackspace.com/information/legal/msa) as it may be updated from time to time.

This Service Description is intended only as a general description of the specified Services and the technical capabilities of Rackspace Technology in relation thereto, which may change from time to time. As applied to Customer, this description may not reflect the actual services provided by Rackspace Technology. In addition, this Service Description may contain one or more additional and/or non-standard services that are not included in the standard Service offering and that may require payment of additional fees. Customer shall not rely on this Service Description.