Prodapt
Chase Extraordinary

An Efficient Observability Implementation Strategy for Digital Service Providers (DSPs)

Improving productivity by 40% with E2E visibility in cloud-native applications

Credits

Srinivasan Murugesan

Sumit Thakur

As DSPs transition towards microservices architecture and cloud-native applications, traditional monitoring tools have shown several limitations

Difficult to get unified analysis with scattered monitoring tools

Different teams (Dev, Ops, and testing teams) use separate tools to monitor logs, traces, metrics, events, and performances. With this, it is not always possible to monitor all metrics and KPIs of full-stack (application, user, business, infrastructure, network, security).

Challenges in correlation and isolation of problems hamper DSPs to deliver on SLA/SLO

Monitoring tools don't give the option to drill down and correlate issues between infrastructure, application performance, and user behavior. Ops and Dev teams often use logs for debugging and performance optimization. This limits the organization's ability to resolve issues quickly and deliver on SLOs and SLAs.

Inability to scale with modern cloud environments

Static dashboards with human-generated thresholds do not scale or self-adjust to cloud environment. This makes it difficult to monitor cloud-native services when there may be thousands of them deployed on a single VM at a given moment in time.

Inflexible to assist in the resolution of unforeseen events

Conventional monitoring relies on alerting only known problem scenarios. As such, there is no visibility into unknown unknowns. These are kind of unique issues that have never occurred in the past and cannot be discovered via dashboards.

DSPs need to make their digital business more **observable** that are easier to understand, easier to control, and easier to fix.

DSPs need to look beyond traditional monitoring - **Observability** enables DSPs to gain critical insights into the performance of today's complex cloud-native environments

The concept of **observability** is gaining rapid momentum as companies accelerate their digital transformation strategies by building out massive cloud-native environments that are inherently hard to observe and operate due to their dynamic and complex nature.

Gartner defines **observability** as the characteristic of software and systems that allows them to be "**seen**" and allows questions about their behavior to be answered.

Observability provides immediate value to the DSP ecosystem by **gaining critical insights from unified data** that were previously siloed. This enables operators to answer questions that were not possible to answer before.

Key Drivers for DSPs to Embrace Observability

Derive insights from unified visibility across ecosystem	Enables powerful analysis by bringing logs, metrics, events, and traces together at scale in a single stack
Improve end-user satisfaction	By reducing the time to identify issues, the improved application uptime and performance will reduce customer churn and enhance return rates.
Lower infrastructure costs	By looking at the data generated, it is possible to optimize infrastructure.

This insight details an **efficient observability implementation strategy** for Digital Service Providers

Efficient Observability implementation strategy for DSPs

While leading DSPs have started implementing observability techniques, many face various implementation challenges and are not able to realize the actual benefit. The **observability implementation strategy** listed here can help DSPs mitigate implementation challenges and gain efficient observability.

1. Build **observability pipeline** based on **OpenTelemetry** standards
2. Promote **observability as a culture** across the organization. Enhance ecosystem observability through **Observability Driven Development (ODD)**
3. Adopt best practices for **data management, security, and governance**
4. Choose the **observability platform** that fits the organization's long-term needs
5. **AI-powered intelligence** for proactive **observability at scale.** Adopt **AIOps strategy.**

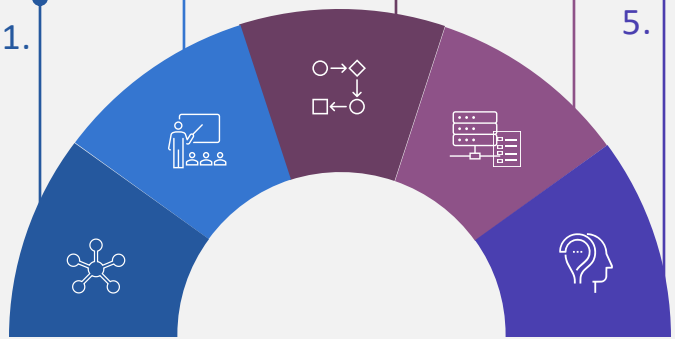

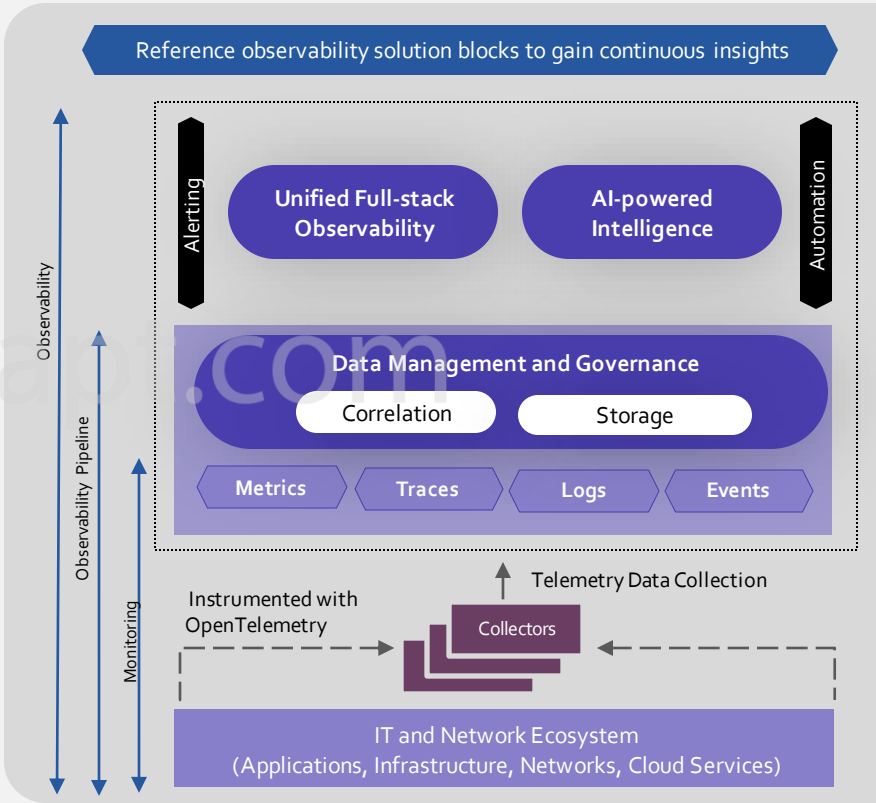

Build Observability pipeline based on OpenTelemetry

Challenges in Observability Pipeline

Lack of standardization of telemetry data leads to:

- Increased complexity to maintain instrumentation (usage of different agents to collect logs, traces, and metrics)
- Vendor lock-in
- Lack of data portability

Tighter coupling of collected data with destination forces teams to use scattered tool sets

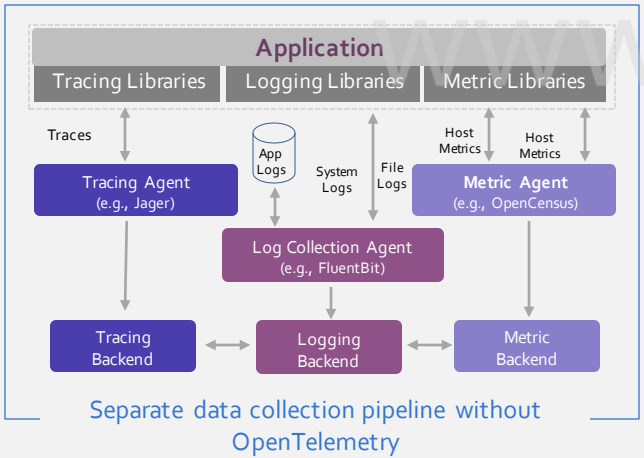

Separate data collection pipeline without OpenTelemetry

Key Recommendations

Build Observability pipeline based on OpenTelemetry standards

With this, it is possible to **decouple the data sources from the destinations** and make the observability data easily consumable.

OpenTelemetry is a Cloud Native Computing Foundation (CNCF) project providing vendor-neutral standards for collecting telemetry data.

OpenTelemetry eases instrumentation for DSPs

- Provides a **single, vendor-agnostic instrumentation** library per language and supports automatic and manual instrumentation.
- Ensures full control of data with the ability to send data to multiple destinations in parallel through configuration.
- Provides a de-facto standard for adding observability to cloud-native applications.
- As OpenTelemetry gains adoption, more frameworks will include out-of-the-box instrumentation.

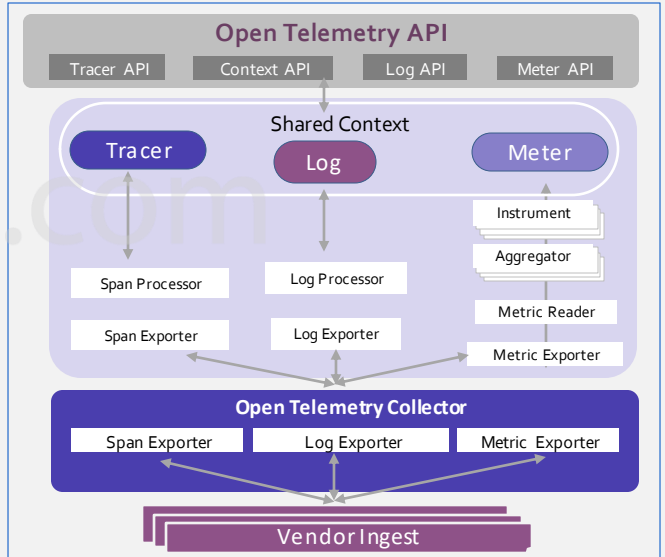

Unified data collection based on OpenTelemetry

Key recommendations for building observability pipeline

Monitor metrics that really matter

- Start with monitoring key metrics that have direct implications on operations and business.
- Establish the baseline list of metrics and optimize it based on the observability learnings.
- To eliminate any capacity issue, it is crucial to focus only on the data sources that hold real value.

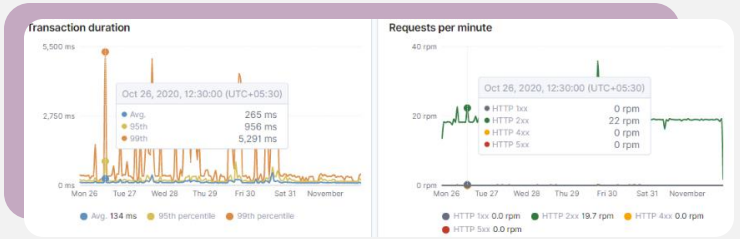

Ensure standard and structured log management techniques

- Structured logging is key for analytics, search, identifying patterns, and improving overall observability.
- **Clearly define logging guidelines**. This should cover key parameters such as log time, log name, format, and content of logs (should comprise of the Correlation ID, Flow ID, Event ID, and Transaction ID).
- Log critical data that helps DSPs to troubleshoot performance problems, solve user experience issues, or monitor security-related events.

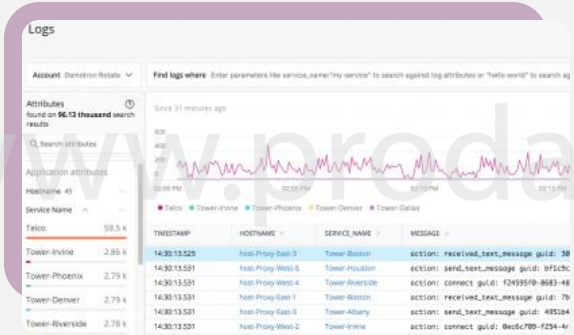

Capture trace IDs to clearly visualize request flow

This enables DSPs to see how a request flows through the system, irrespective of whether you're using a service mesh, or even a load balancer or proxy.

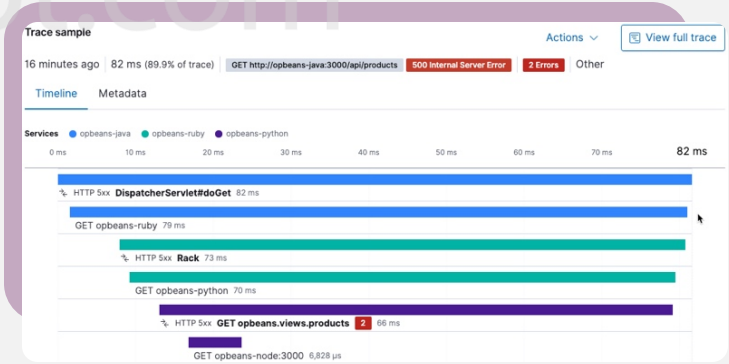

Make log level configurable

- This helps to adjust the verbosity of logs and get enough information as needed.
- **Messaging queue topics should be retained for at least a week** for reliable debugging.

Promote observability as a culture across the organization

Enhance ecosystem observability through Observability Driven Development (ODD)

Challenges

In traditional monitoring, visibility is **not a consideration during the design or development phase**. As such, the DevOps team is aware of issues only when services fail or are about to fail in predictable ways.

As microservices, containers, and other new architectural components make systems more distributed, complicated, and unpredictable, there is an **increase of unknown unknowns.**

Incorporate observability aspects throughout the software development life cycle by **ensuring Observability Driven Development (ODD)**

Key Recommendations

Promote **observability as a culture** across the organization

- Observability as a culture is the degree to which a team or company values the ability to inspect and understand systems, their workload, and their behavior.
- Observability Driven Development (ODD) enables observability culture and tackling the unknowns.

- Determine what to measure based on QoS and KPIs to be met.
- Identify appropriate places where instrumentation needs to be added.

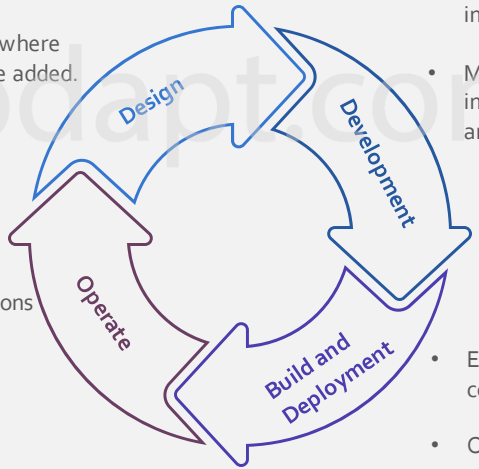

- Standardize the context and have sufficient context included consistently across all instrumentation data.
- Maintain a right balance on the level of instrumentation, or else this can overwhelm analysis.

- Feedback of observability learnings from the operations team to the development team is essential for continuous improvement.
- Proactive monitoring and alerting should be in place.

- Enforce observability as part of the continuous deployment process.
- Observe for unusual behavior at an early stage through automation.

Adopt best practices for data management, security, and governance

Challenges

Storage capacity issues

- **Overlogging leads** to a situation where log storage capacity is consumed quickly.
- **Lack of retention policies** leads to quick exhaustion of storage capacity leading to cost increase and operational issues.

Security compliance

- Lack of role-based access and GDPR non-compliance could lead to severe security breaches and penalization.

Key Recommendations

Don't do analysis in silos Centralize and correlate all data sources

- A **single pane of view** helps to connect dots between captured logs, events, traces, and metrices. This gives the whole story of what's happening at any point in time.
- Logs from disparate sources (such as application, system, server, and database) can be collected, parsed, and **stored in a central repository with indexing.** This enables analytics to be more efficient.

Create a flexible data retention policy

- Clearly **define the duration of retention** for various types of data (e.g., regulatory data, machine state data, etc.)
- **Follow the 3-2-1 rule for storage and backup** – There should ideally be 3 copies of the data, stored on 2 different media, with at least 1 stored off-site or in the cloud.
- **Log storage must work as a cyclic buffer** – This deletes the oldest data first when the storage limit is reached.

Implement security policies for collected data

- **Role based access control** should be implemented for access to storage data.
- Make sure sensitive data get **anonymized or encrypted.** For e.g., collected logs might have personal data that must be encrypted to avoid GDPR breach.

Use stored data (logs) to identify automation opportunities

- Logging should be viewed as an **enabler for automation (DevOps and CI/CD)** in addition to just troubleshooting. Capture **where** the issues are introduced and **what** are the sources of these issues to identify an **automated fix.**

Choose the observability platform that fits the organization's long-term needs

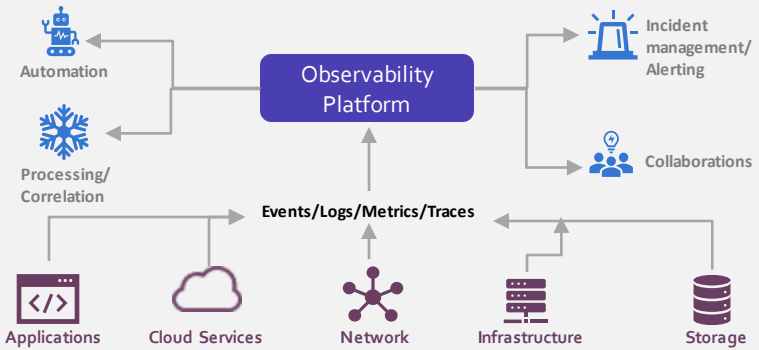

DSP can select the right observability platform based on the following parameters

- Full stack monitoring (cloud, business, user, applications, infrastructure, network)
- Open source (e.g., SolarWinds, Elastic) or commercial tool (e.g., AppDynamics, Dynatrace, etc.)
- Scalable and fits the organization's long-term strategy
- Supports OpenTelemetry (e.g., Elastic, NewRelic, etc.)
- Supports intelligence and AIOps (e.g., Elastic, Dynatrace, AppDynamics, MoogSoft, etc.)
- Ability to correlate metrics, traces, logs, and events to business outcomes
- Real-time analysis (aggregation and visualization)

Recommendations for effective implementation of observability platform

Use distributed tracing capability to get a clear view of granular insights

- Understand how the services are interacting.
- Visualize service calls between systems; find where latency issues are arising in the end-to-end flow.
- Pinpoint the components that need optimization.

Use debugging strategies such as:

- Comparing changes across time series to detect performance degradation.
- Deriving insights from user behavior patterns.
- Analysing impacts based on the context of errors/alerts.

Apply alerting techniques

- Start with foundation alerts and over a period, optimize alerting strategy.
- Fewer alerts focussing on a core set of metrics is key.

Single platform for E2E observability

DSPs can move away from the usage of scattered tools (such as Prometheus for Metrics, Grafana for custom dashboards, and Jaeger for Traces) towards a single observability platform.

AI-powered intelligence for proactive observability at scale. Adopt AIOps strategy.

Challenges

Inability to scale

As the scale increases dramatically with the digitization of business

- Traditional tools aren't designed for this scale
- Reactive monitoring hurts performance, customer experience, and business

Difficult to cut through the noise to identify root causes of performance issues with traditional tools

Use machine learning models and AI for sophisticated data analysis and to provide actionable answers to problems

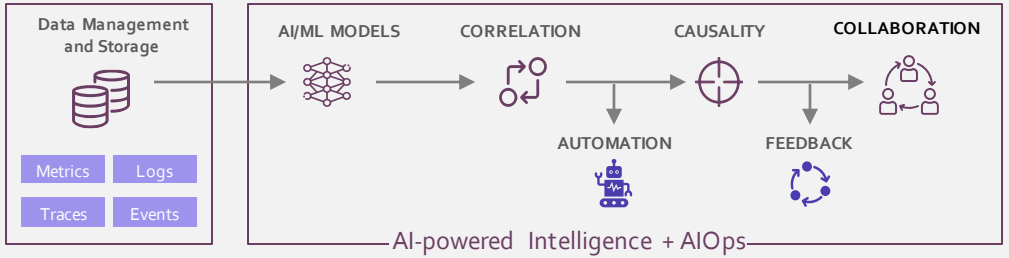

AIOps is becoming an embedded capability of observability.

Gartner predicts that exclusive use of AIOps and digital experience monitoring tools to monitor applications and infrastructure will rise from **5% in 2018 to 30% in 2023**

Recommendations

AI-powered observatory tool

- Enterprises with non-AI-based observatory tool can either switch to an AI-powered observatory tool or develop a custom AI–powered layer on existing APM to move towards proactive observability at scale.

Adopt AIOps strategy

- **Intelligent alerting** to indicate an emerging issue. Filter and correlate data into incidents to avoid alert fatigue. Separate low entropy events, which could end up as false alarms.
- **Automated remediation** of common issues. Common issues could be known issues or unknown issues. E.g., Restarting a component or cleaning up a full disk can be handled automatically.

Use cases are emerging

- Automate root cause analysis and assess the impacts to the business by aggregation of data and creation of causality/relationships.
- Predict fault, failure, forecast, capacity based on trend analysis.
- Auto instrumentation of data
- Auto recommendation of key KPIs and metrics

Enterprises need to prioritize the development of an AIOps strategy. AI-powered observatory tool combined with AIOps strategy for observability at scale can simply the demands of increasingly complex ecosystem.

Sample use case for AI-powered observability– Addressing scalability issue for a leading DSP in Europe using ML-based capacity modeling for video streaming ecosystem

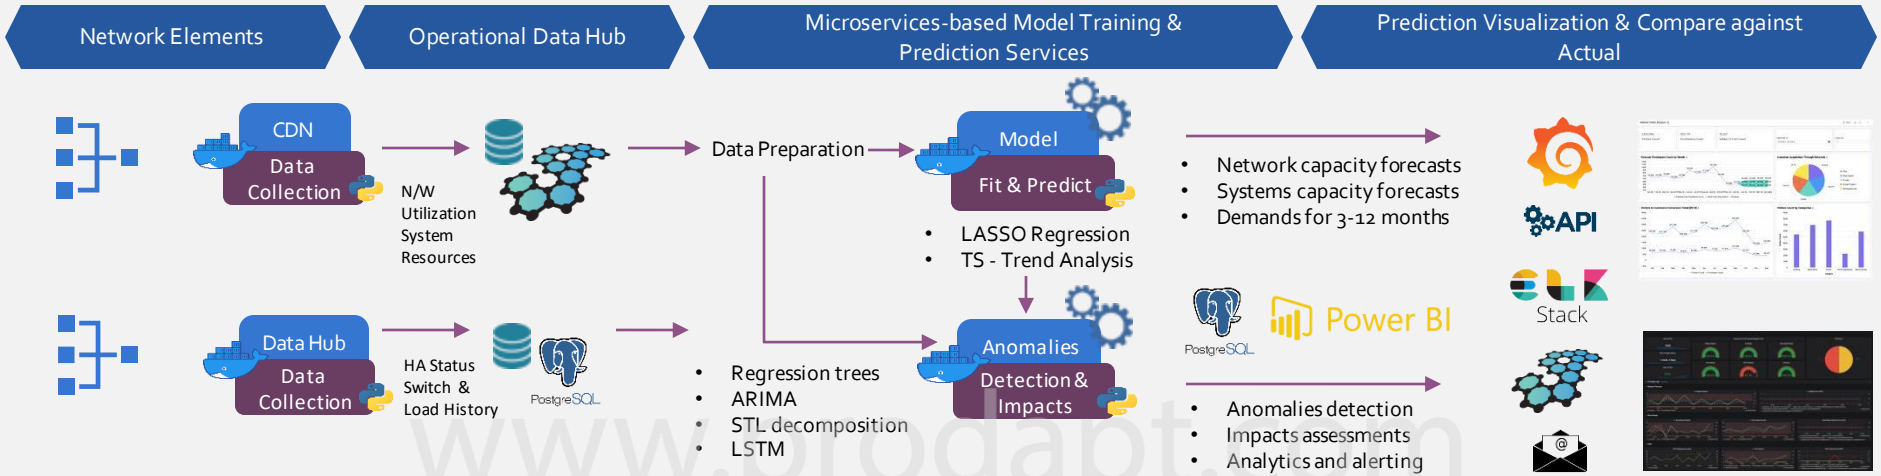

Key Steps

- Data is ingested from Big Data system using HIVE
- Stages of prediction service is deployed as microservices & dockers
- ML activity involves capacity forecasting using a combination of **regression-based trend (LASSO),** time-series, and Gaussian noise
- Capacity metrics are forecasted for the next 3, 6, 9, 12 months
- All forecasts/prediction results are pushed to ELK/Grafana for visualization

Business Value Delivered

Zero capacity failures observed in past one year for the DSP having its presence in **15+** countries with more than **25M** customer base.

New areas of **capacity risks in different parts** of the network were **identified much in advance.**

Enabled DSP to timely accommodate and plan for the required future capacity.

80% accurate prediction of CDN and back-office capacity using machine learning.

Business benefits achieved by a leading DSP in Europe after observability implementation using the proposed strategy

Get in touch

USA

Prodapt North America, Inc.
Oregon: 10260 SW Greenburg Road, Portland
Phone: +1 503 636 3737

Dallas: 1333, Corporate Dr., Suite 101, Irving
Phone: +1 972 201 9009

New York: 1 Bridge Street, Irvington
Phone: +1 646 403 8161

CANADA

Prodapt Canada, Inc.
Vancouver: 777, Hornby Street,
Suite 600, BC V6Z 1S4
Phone: +1 503 210 0107

PANAMA

Prodapt Panama, Inc.
Panama Pacifico: Suite No 206, Building 3815
Phone: +1 503 636 3737

UK

Prodapt (UK) Limited
Reading: Suite 277, 200 Brook Drive,
Green Park, RG2 6UB
Phone: +44 (0) 11 8900 1068

IRELAND

Prodapt Ireland Limited
Dublin: Suite 3, One earlsfort centre,
lower hatch street
Phone: +44 (0) 11 8900 1068

EUROPE

Prodapt Solutions Europe &
Prodapt Consulting B.V.
Rijswijk: De Bruyn Kopsstraat 14
Phone: +31 (0) 70 4140722

Prodapt Germany GmbH
Münich: Brienner Straße 12, 80333
Phone: +31 (0) 70 4140722

Prodapt Digital Solution LLC
Zagreb: Grand Centar,
Hektorovićeva ulica 2, 10 000

Prodapt Switzerland GmbH
Zurich: Muhlebachstrasse 54,
8008 Zürich

Prodapt Austria GmbH
Vienna: Karlsplatz 3/19 1010
Phone: +31 (0) 70 4140722

SOUTH AFRICA

Prodapt SA (Pty) Ltd.
Johannesburg: No. 3, 3rd Avenue,
Rivonia
Phone: +27 (0) 11 259 4000

INDIA

Prodapt Solutions Pvt. Ltd.
Chennai: Prince Infocity II, OMR
Phone: +91 44 4903 3000

"Chennai One" SEZ, Thoraipakkam
Phone: +91 44 4230 2300

IIT Madras Research Park II,
3rd floor, Kanagam Road, Taramani
Phone: +91 44 4903 3020

Bangalore: "CareerNet Campus"
2nd floor, No. 53, Devarabisana Halli,
Phone: +91 80 4655 7008

THANK YOU!