

# SEVONE DATA CLOUD

## THE MONITORING & AIOPS PLATFORM FOR NETWORK STREAMING TELEMETRY

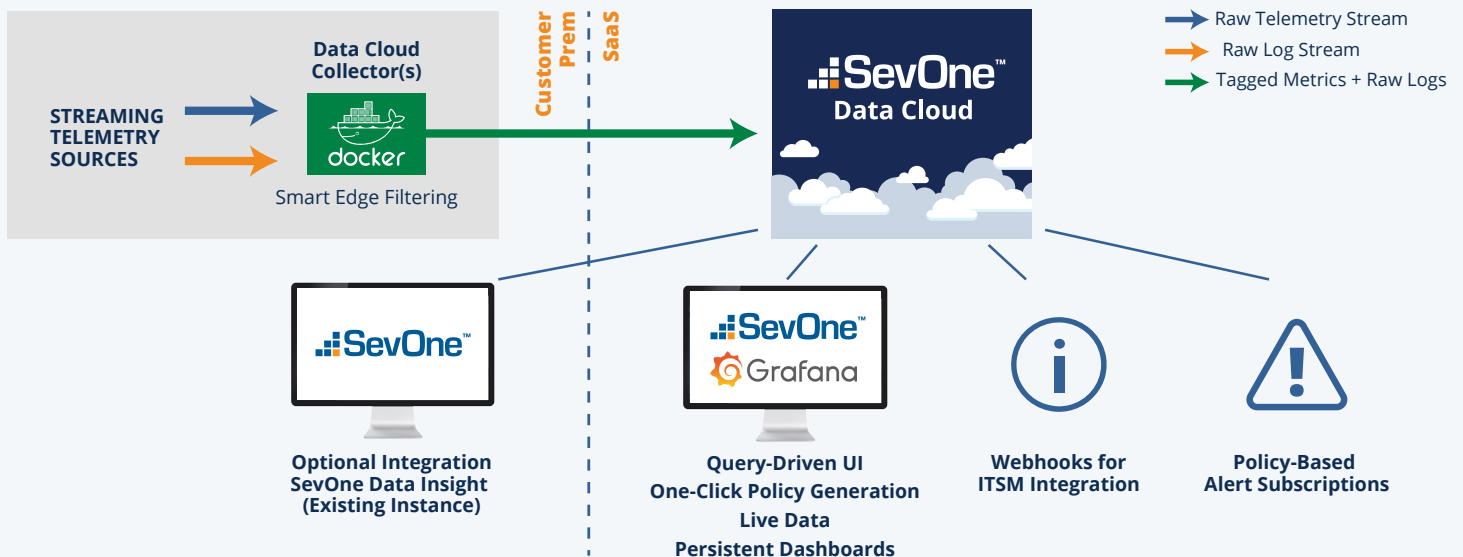
There is a new reality. Applications and networks no longer just support your business—they are your business. In an always-on, always-connected digital world, the expectation is for enterprises, carriers and managed service providers to deliver flawless network and application experiences to their employees, customers and business partners—every second of every day.

As networks transition to support the next generation cloud-based applications and services driving digital transformation, more granular and deterministic performance data is critical for continuous operations. Legacy management systems depending on once every five-minute, poll-based metrics from across your infrastructure may limit the efficiency and reliability of networks and applications and the overall businesses they support.

To address this challenge, leading network equipment providers are beginning to support a “push” type of model to provide performance data eliminating the need for management systems to poll these new devices. Now management systems can request specific telemetry data be sent to them at specific intervals—intervals as fast as every second, but that can lead to up to 300x more data to be collected and analyzed in real-time.

SevOne Data Cloud, the Monitoring & AIOps Platform for Network Streaming Telemetry, currently in open trial, is ready to address this market challenge with:

- Real-time and historical insight of network streaming telemetry and syslogs at cloud scale
- Unlimited service-based tagging for business context
- Real-time and on-demand policy-based analytics
- A powerful, fully-managed SaaS offering



### STREAM.

Subscribe to per second streams of network telemetry

### TAG.

Converts raw performance data & logs to tagged-based metrics

### ANALYZE.

Policy-based thresholding, baselining, aggregation and alert subscriptions

## KEY FEATURES

### Real-time & Historical Insight of Network Streaming Telemetry & Syslogs at Cloud Scale

- **Network Streaming Telemetry Ingestion:**  
Transform raw telemetry into stream of metrics  
Smart Edge filtering – analyze the data that matters
- **Syslog Ingestion**  
Raw syslog collection  
Transform raw syslog into stream of log metrics
- **Live Dashboards**  
Visualize streaming network telemetry and syslog metrics with persistent, live, Grafana-based dashboards  
Integrate with an existing SevOne Data Insight implementation
- **Query-based UI:**  
Ad-hoc queries of real-time and historical metrics  
High cardinality resource selection

### Unlimited Service-Based Tagging for Business Context

- **Default Tagging**  
Raw data is tagged by the collector based on known data structure from the source device
- **Flexible User-Based Tagging**  
Users can add their own tagging to meet their business need/vocabulary  
Unlimited tags can be added by users—against real-time and historical metrics
- **Leveraged System Wide**  
Tagging is used in searches, aggregations of indicators, policy creation, data queries, event stream subscriptions, dashboard creation and more
- **SevOne Data Model for AIOps – Key to AIOps Success**  
New data model with separate data stores for “metric/value” and “tags”  
Dynamic association of data stores fuels on the fly grouping, aggregations and policy creation

### Real-Time and On-Demand Policy-Based Analytics

- **Real-Time Analytics**  
Real-time analytic policies set by users run continuously against real-time streams at ingestion—including, aggregations, baselines and thresholds
- **On-Demand, Run Time Analytics**  
Leverage more than a dozen different analytics , including transformation, smoothing, bounds and seasonality-based analytics can be requested to run against queried data
- **One-Click Policy Generation**  
Create policies based on historical metrics  
Initiate real-time analytics of new, incoming streams of metrics
- **Event Policy Subscriptions**  
Users can subscribe to tag-based policy events  
Integrate with ITSM systems via webhooks  
First occurrence and persistent event streams

### In a Powerful SaaS Offering

- **Fully managed offering**
- **100% SaaS**
- **Now in open trial**

## KEY BENEFITS

SevOne Data Cloud enables network and IT operations, network engineering and data science teams to achieve their business agility, efficiency and reliability goals with:

- **Faster, more accurate troubleshooting and SLA compliance**
- **Granular policy-based event detection and alert subscriptions to drive automation**
- **More precise capacity planning and optimization**
- **New business value discovery through flexible data science analytics**

## NETWORK STREAMING TELEMETRY

SevOne Data Cloud supports the collection, tagging, policy-based analysis and troubleshooting of network streaming telemetry and syslog data.

Our support for network streaming telemetry includes:



- Cisco IOS XR
- Cisco IOS XE
- Supported Protocols: 32 bit TCP, 64 bit gRPC

## SYSLOG RFCs

- 3164, 5424, 5425 & 6587

If you have additional device support needs let us know, we'd be happy to work with you.

## ADVANCED ANALYTICS:

SevOne Data Cloud enables users to query all stored metrics and run a range of advanced analytics against that query—then in one click turn that query into an automated policy that is run continuously against any new incoming streaming data.

DATA CLOUD ANALYTICS	FEATURE	BENEFIT
<b>ON DEMAND ANALYTICS: BOUNDS</b>		
<b>Exponential Standard Deviation</b>	<p>Draws standard deviation bounds above &amp; below lines on a baseline statistic.</p> <p>Exponential standard deviation is a measurement of variation among values in a data set, giving greater weight to more recent data points.</p>	Exponential standard deviation bounds provide a visual aid showing the expected range of many data types, with an emphasis of the range of recent data.
<b>Group Percentile</b>	<p>Draws upper and lower bounds containing the specified percentile of data within the matched group of indicators.</p> <p>For example, a value of 95 will produce bands containing 95% of the data.</p>	Group percentile bounds provide a visual aid showing the range of the matched indicators at each point in time, without being biased by extremely large or small values.
<b>Group Standard Deviation</b>	<p>Draws standard deviation bounds above &amp; below a baseline statistic (commonly the Group Average).</p>	Group standard deviation bounds provide a visual aid showing the expected range of the matched indicators at each point in time. These bounds are appropriate for many data types.
<b>Percentile</b>	<p>Draws upper and lower bounds containing the specified percentile of data within the data time window. e.g. a value of 95 will produce bands containing 95% of the data.</p>	Percentile bounds provide a visual aid showing the range of data values, without being biased by extremely large or small values.
<b>Standard Deviation</b>	<p>Draws standard deviation bounds above &amp; below lines on a baseline statistic.</p> <p>Standard deviation is a measurement of variation among values in a data set.</p>	Standard deviation bounds provide a visual aid showing the expected range of data, and are appropriate for many data types.
<b>ON DEMAND ANALYTICS: GROUP SMOOTHING PRESET</b>		
<b>Group Average</b>	<p>Overlay the mean of all matched indicators at each 1 minute time bucket.</p> <p>Group Average overlays the mean of matched indicators at each point in time.</p>	This preset enables the user to identify outliers at each point in time among a group of indicators.
<b>Group Median</b>	<p>Overlay the median of all matched indicators at each 1 minute time bucket.</p> <p>Group Median overlays the median of matched indicators at each point in time.</p>	This preset enables the user to identify outliers at each point in time among a group of indicators, without being biased by extremely large or small values.
<b>ON DEMAND ANALYTICS: GROUP TRANSFORMATION PRESET</b>		
<b>Group Rank</b>	<p>Order matched indicators within each 1 minute time bucket.</p>	This preset displays the relative position of matched indicators at each point in time.

DATA CLOUD ANALYTICS	FEATURE	BENEFIT
<b>ON DEMAND ANALYTICS: SMOOTHING PRESET</b>		
<b>30 Minute Moving Average</b>	<p>Smooth the data using the average (mean) over a 30 minute moving window.</p> <p>Moving Average overlays the indicators mean to serve as a baseline. The mean is equally sensitive to old data and new data in the time window.</p>	<p>A short-term moving average to help the user judge whether the current data point is departing from its recent operating range.</p>
<b>Hourly Moving Median</b>	<p>Smooth the data using the median over a 1 hour moving window.</p> <p>Moving Median overlays the indicator's median over the statistic time window.</p>	<p>The moving median baseline is useful for data containing large outliers which would bias a moving average computation. This helps the user judge whether the current data point is within expectations.</p>
<b>Weighted Daily Seasonality</b>	<p>Calculate a baseline for data having daily seasonality, using an exponential moving average truncated to 2 weeks of historical data.</p> <p>Exponential moving averages place more weight on recent data than older data. Seasonal sampling is used to compute the average from previous samples taken at the same time of day.</p>	<p>This baseline preset is appropriate for data having daily seasonality. It computes and displays the expected data range with an emphasis on recent days, and helps the user judge whether the current data point is within expectations.</p>
<b>Weighted Weekly Seasonality</b>	<p>Calculate a baseline for data having weekly seasonality, using an exponential moving average truncated to 3 months of historical data.</p> <p>Exponential moving averages place more weight on recent data than older data. Seasonal sampling is used to compute the average from previous samples taken at the same time of day and day of week.</p>	<p>This baseline preset is appropriate for data having weekly seasonality. It computes and displays the expected data range with an emphasis on recent weeks, and helps the user judge whether the current data point is within expectations.</p>
<b>ON DEMAND ANALYTICS: TRANSFORMATION PRESET</b>		
<b>30 Minute Standardization</b>	<p>Standardize indicator data by the mean and standard deviation on a 30 minute moving window.</p> <p>Standardize scales each data point by subtracting its moving average and dividing by its standard deviation.</p>	<p>This preset places values for all matching indicators on a common scale so they can be compared with each other.</p>
<b>ON DEMAND ANALYTICS: GRANULARITY</b>		
<b>Raw, Olympic Average, Average, Total, Min &amp; Max</b>	<p>Raw = as collected</p> <p>Olympic Average = Average of raw data, after removing the highest value and the lowest value.</p> <p>Average = Mean of raw data values</p> <p>Total = Sum of raw data values</p> <p>Min= Lowest value of selected data set</p> <p>Max = Highest value of selected data set</p>	<p>When a granularity other than Raw is selected, data is aggregated into non-overlapping buckets with a fixed time duration. This uniform spacing is important for computing seasonal and group statistics.</p>

DATA CLOUD ANALYTICS	FEATURE	BENEFIT
<b>REAL-TIME ANALYTICS</b>		
<b>Event Streams</b> <ul style="list-style-type: none"> <li>■ Static Threshold <ul style="list-style-type: none"> <li>■ Raw</li> </ul> </li> <li>■ Baseline Comparison (Multiple or Offset) <ul style="list-style-type: none"> <li>■ Raw</li> </ul> </li> </ul>	<p>A tag query is used to select indicators as members of the policy. New values for members of the policy are ingested by Data Cloud and compared against thresholds set up in the policy. The threshold can also be a comparison of the ingested value to a baseline of the indicator over time.</p> <p>Indicators that break the thresholds trigger event streams. Once a threshold is no longer breached, the event streams are closed.</p> <p>Event streams are permanent and will re-open if a threshold is triggered in the future.</p>	<p>Users can enable continuous analysis of specific combinations of new streaming metrics at ingestion, allowing them to be notified of baseline alert conditions to meet their business needs.</p>
<b>Time Over Threshold</b> <ul style="list-style-type: none"> <li>■ Continuous Time Over Threshold</li> </ul>	<p>Continuous Time Over Thresholds behaves much like Event Streams described above, except that indicator values are compared to thresholds over various factors of time.</p> <p>Continuous Time Over Threshold requires that an indicator continuously breach a threshold for the configured duration before an event is triggered. Example - 1 minute continuously triggering time over a 5 minute window.</p>	<p>Users can enable continuous analysis of specific combinations of new streaming metrics at ingestion, allowing them to be notified of baseline alert conditions across a user-define time period to meet their business needs.</p>
<b>Aggregations</b>	<p>A tag query is used to select indicators as members of the policy. New values for members of the policy are streamed into SevOne Data Cloud for analysis. All available values are rolled up into a 1 minute bucket containing the following summary statistics:</p> <ul style="list-style-type: none"> <li>• Sum</li> <li>• Count</li> <li>• Min</li> <li>• Max</li> <li>• Average</li> </ul>	<p>Users can create aggregation policies to, for example, summarize all production interfaces within their business regions. Users could then monitor those business region-based aggregates for anomalies.</p>

GET STARTED WITH SEVONE:  +1.302.261.8718  solutions@sevone.com  www.sevone.com