APPLICATION NOTE

Intel Corporation
Software Defined Datacenter Solutions Group

Collecting & Monitoring Platform & Container Telemetry

Introduction

When hardware or software fails on the backend, it needs to be fixed immediately, and preferably automatically, to prevent end users from experiencing painful service disruptions. As the industry accelerates on its path to cloud-native, drawn by the promise of rapid innovation, fast-time-to-market and lower total cost of ownership—developers and architects must consider how to monitor and report on system performance within existing (brownfield) deployments, while simultaneously laying the groundwork for future (greenfield) deployments.

The migration to a fully cloud-based model won't happen overnight, and it's essential that robust service assurance capabilities are maintained as the transformation to a software-defined and increasingly virtualized network environment continues. Systems must be vigilantly monitored for utilization and malfunctions to mitigate (and rapidly recover from) disruptions, using both platform and container telemetry. Telemetry information serves as a vital resource for analyzing the platform's state and health, and for delivering on service assurance goals.

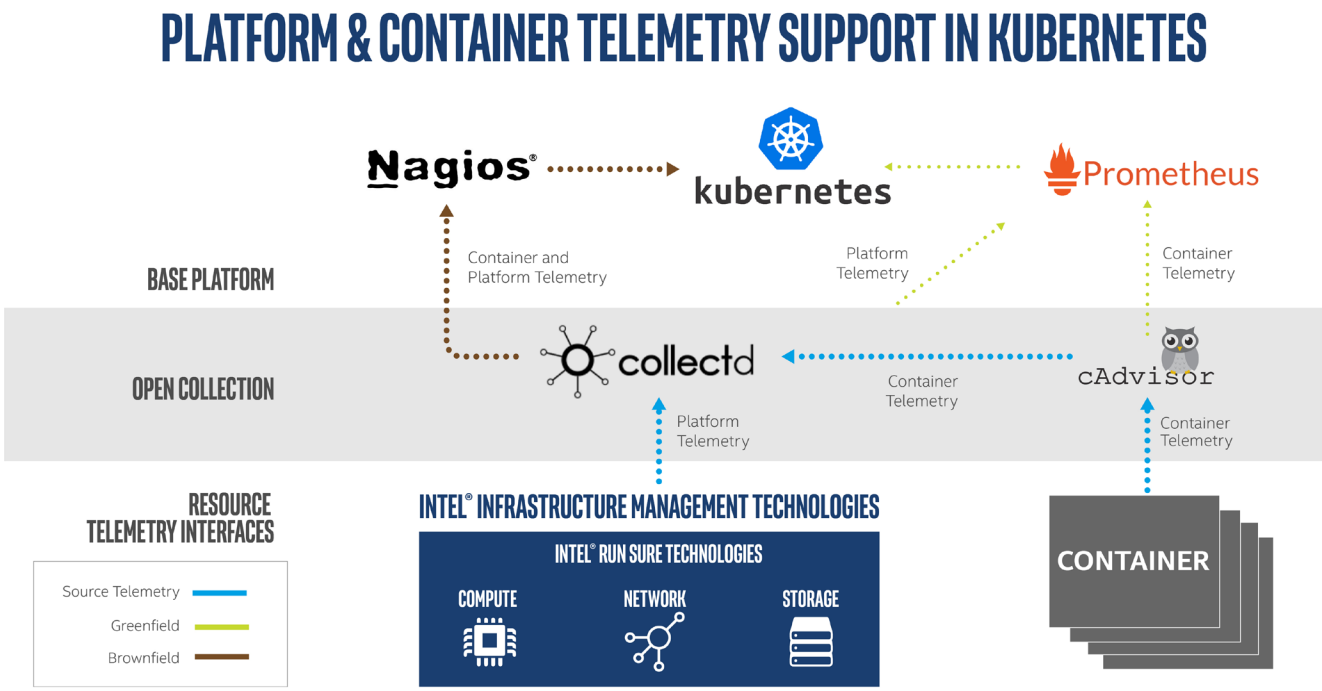

Interoperability between today's and tomorrow's environments is a critical component of facilitating network transformation. This application note explores how developers and architects can bridge the gap, integrating platform and container telemetry with existing monitoring solutions in brownfield environments. It outlines how collectd, the industry's collection agent of choice, can integrate container telemetry with an existing FCAPS (Fault, Configuration, Accounting, Performance, Security) management system that uses SNMP (Simple Network Management Protocol) in order to gather the same platform telemetry from containers as is available from VMs.

The document also demonstrates how platform telemetry—integrated with a monitoring solution like Nagios—can remediate platform failures where containers are in operation, highlighting the value of bringing these components together. It provides two use cases, utilizing a small subsection of the many platform and container telemetry sources available, to bring these capabilities to life. The first example uses collectd to feed container telemetry to a SNMP collector. The second makes use of Nagios open source (and other) software elements, gathering container and platform telemetry to monitor the infrastructure and provide fast corrective action.

Both discrete use cases subscribe to platform reporting features, such as Intel® Run Sure technologies and Intel® Infrastructure Management Technologies. These metrics or telemetry are plumbed into the container orchestration engine via several open source components, primarily collectd, cAdvisor, and, in the case of kubernetes, Prometheus. They illustrate how telemetry can serve as the source for advanced provisioning and service assurance, enabling not only corrective but also preventative action as a positive outcome of capacity management.

Telemetry Use Cases

Monitoring with Container Telemetry & SNMP

One method developers and architects have at their disposal to monitor and report on container metrics and events in varied enterprise and telco environments is by utilizing SNMP, a protocol used by existing FCAPS management systems. This use case leverages previously deployed monitoring mechanisms, specifically collectd with SNMP, to support both virtualized environments (virtual machines and containers) and brownfield deployments.

While they are configured for virtualized network functions in this example, these daemons can also support the coming cloud-native applications. The use case illustrated below leverages Telegraf and InfluxDB as an SNMP collector, and Grafana to visualize the collected metrics using SNMP.

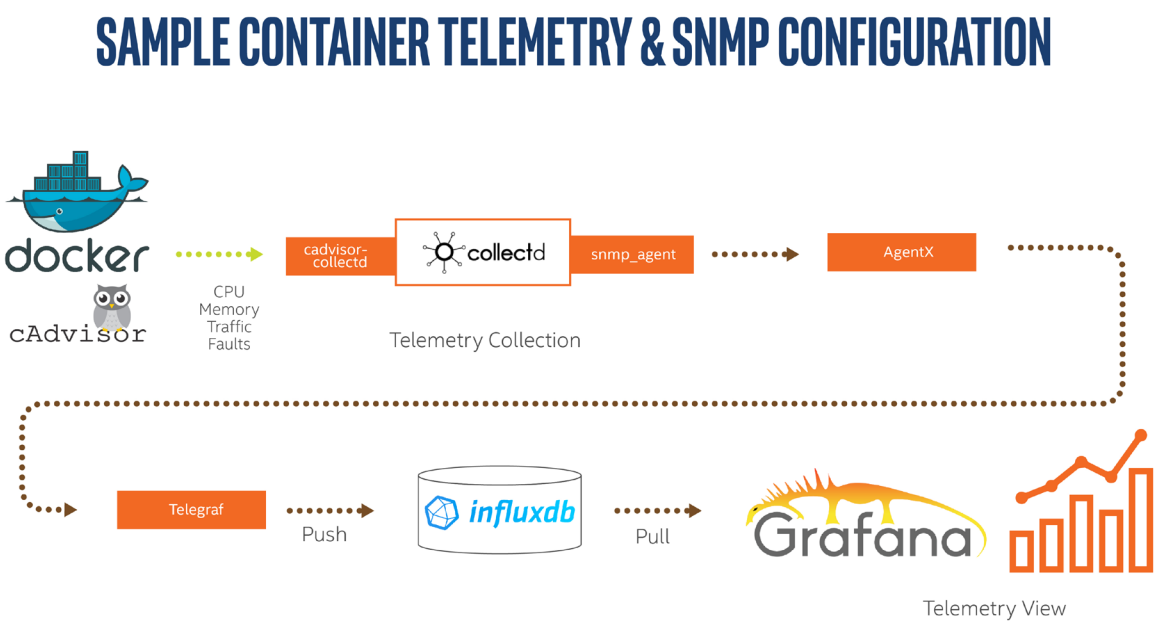

In this dataflow, collectd serves as the broker that brings the components together, using the cAdvisor-collectd plugin to gather data on statistics or events, before publishing it over SNMP in reply to GET, GET_NEXT or WALK requests issued periodically from Telegraf. The SNMP payloads are then captured by Telegraf and pushed into the Influxdb time series database via an HTTP interface. Finally, the model or times series data is presented to the user using user-defined graphs in the Grafana dashboard. Grafana periodically queries Influxdb and updates the on-screen graphs.

Sample Container Telemetry & SNMP Elements

Software Component	Description
cAdvisor	A container advisor that provides resource usage and characteristics information on running containers
collectd	A statistics daemon that collects and reports on system and application performance metrics
cAdvisor-collectd	A metrics plugin that queries cAdvisor periodically for container metrics
snmp_agent	An agent that reports collectd metrics over an SNMP interface to Telegraf using AgentX
AgentX	A networking protocol that communicates with objects as defined by master agents and subagents in the SNMP protocol
Telegraf	An SNMP agency that collects and reports metrics into InfluxData (Influxdb) over the HTTP
Influxdb	An open-source, scalable time-series database that stores metrics, events, and performance analytics
Grafana	A front-end platform that visualizes time series data, typically used in conjunction with a time series database such as Influxdb.

Monitoring with Platform/Container Telemetry & Nagios

Another method developers and architects can now leverage is Nagios. Nagios, an open source tool historically used in IT infrastructure, can now be used in Kubernetes containers as well, not just for monitoring but also to take corrective action. Intel® Infrastructure Management Technologies are monitored for both corrected and uncorrected errors using platform telemetry data. These errors are reported to Nagios, which is configured to drain a Kubernetes node of its containers when a certain threshold of uncorrected errors occurs.

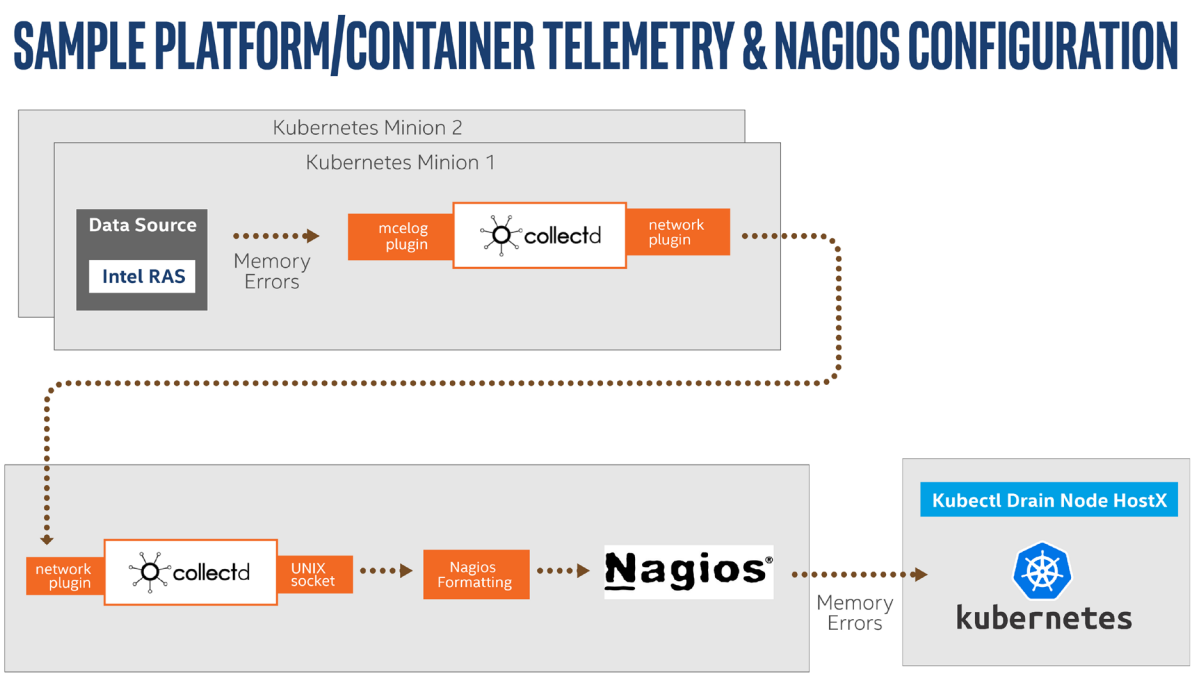

When Nagios receives an event about an uncorrected memory error (indicating a DIMM failure on the platform), Nagios executes a script like the one illustrated in the example above, resulting in Kubernetes draining the target node of containers. The Kubernetes scheduler then starts up the containers on other available nodes, preventing service disruptions.

collectd is central to this dataflow as well. In this example, MCE-inject is used to inject simulated memory errors used for smoke tests; in that it allows user to generate events, as a means to test the reporting and corrective actions. mcelog is the primary collection component. It's responsible for capturing the hardware events that need to be acted upon.

collectd retrieves those events from mcelog using an mcelog plugin, then publishes these metrics over the network to the Nagios server with the collectd network plugin. Once the payload is received, an event handler in Nagios takes corrective action. In this case, a script is executed, causing Kubernetes to drain the impacted server of containers on the host.

Sample Platform/Container Telemetry & Nagios Elements

Software Component	Description
mcelog	A hardware daemon that accounts for and logs RAS machine check exceptions in platform subsystems, such as memory, IO, QPI, and CPU
collectd	A statistics daemon that collects and reports on system and application performance metrics
collectd-mcelog	A plugin that watches the syslog (and other agents) for mcelog reports and publishes them via collectd
collectd_nagios	A plugin that reports collectd events to Nagios

Summary

Intel's Data Center Group ensures that enterprises, cloud and communication service providers, as well as the developers and architects that support them, have the tools and education they need to transform their network environments, while still maintaining Quality-of-Service (QoS). Through the development of Intel® Run Sure* and Infrastructure Management Technologies**, the platform telemetry they provide, and our contributions to collectd, the industry's agent-of-choice for the collection of both platform and container telemetry—we deliver solutions that can not only support brownfield deployments, but also the coming wave of cloud-native applications.

As hardware and software become progressively disaggregated, and services are pushed out more dynamically—complexity will inevitably increase. Employing proven solutions, designed to work across a range of environments can reduce this complexity, easing the transition to a cloud-based model, while mitigating risks. This paper explored two paths that developers and architects can readily employ to monitor usage and potential malfunctions within container environments, using collectd to gather platform and container telemetry.

Interoperability between today's and tomorrow's environments is a critical component of facilitating network transformation. Intel's Data Center Group is at the center of this transformation, providing valuable knowledge and tested solutions to meet existing and emerging customer needs.

Additional Resources:

*Intel Network Builders: NFV Service Assurance

**Feature Brief: Reliability, Availability, & Serviceability (RAS) of Intel® Infrastructure Management Technologies Feature Support

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software, or service activation. Learn more at intel.com, or from the OEM or retailer.

No computer system can be absolutely secure. Intel does not assume any liability for lost or stolen data or systems or any damages resulting from such losses.

You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document. The products described may contain design defects or errors known as errata, which may cause the product to deviate from published specifications. Current characterized errata are available on request.

This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest Intel product specifications and roadmaps.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade. Intel® Turbo Boost Technology requires a PC with a processor with Intel Turbo Boost Technology capability. Intel Turbo Boost Technology performance varies depending on hardware, software and overall system configuration. Check with your PC manufacturer on whether your system delivers Intel Turbo Boost Technology. For more information, see www.intel.com/technology/turboboost.

Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Consult other sources of information to evaluate performance as you consider your purchase. For more complete information about performance and benchmark results, visit www.intel.com/performance. Results have been estimated or simulated using internal Intel analysis or architecture simulation or modeling, and provided to you for informational purposes. Any differences in your system hardware, software, or configuration may affect your actual performance.

Copies of documents that have an order number and are referenced in this document may be obtained by calling 1-800-548-4725 or by visiting www.intel.com/design/literature.htm.

Intel, the Intel logo, and Intel® Xeon® are trademarks of Intel Corporation in the United States. and/or other countries.

*Other names and brands may be claimed as the property of others.

 12/18/HM/PMG/PDF001

Jan 2018 SKU 337045-001 Collecting and Monitoring Platform and Container Telemetry - Application Note