

Optimize the Performance of Your Cloud Infrastructure

AppFormix software leverages cutting-edge Intel Resource Director Technology (RDT) hardware features to improve cloud infrastructure monitoring and control

Challenge


Telecoms and operators of scale-out cloud infrastructure need to know in real time how infrastructure is performing relative to SLAs. Traditional monitoring software takes too long to discover problems and issue alerts, which then require human intervention. Operators want a solution built for the cloud-native, DevOps world.

Solution

AppFormix detects physical and virtual infrastructure issues and automatically manages remedial action based on predefined SLAs. It uses machine learning, thin agents on each node, and a dedicated message bus to automate monitoring for environments where application or network function performance demands the fastest possible response time.

Benefits

- Cloud optimization that works in real time to keep apps and network functions performing up to par
- Purpose-built for DevOps and cloud-native environments
- Delivers excellent performance for both VM and container clouds



In combination, the Juniper Networks® AppFormix orchestration and analytics platform and Intel's Resource Director Technology (RDT) hardware, benefit software-defined infrastructure (SDI) used in cloud and NFV solutions that are the new basis of a myriad of enterprise IT and service provider use cases. This joint solution enables prioritization of important container- or VM-based applications and virtual network functions, using detection of and mitigation of noisy neighbors, and improving the workflow between infrastructure maintainers and software developers in a modern collaborative DevOps environment.

The Challenge

The tools we've used for a generation to optimize infrastructure performance are insufficient when applied to the architectures of a modern cloud. Two foundational shifts are fundamentally redefining how IT operations professionals are preparing to meet the needs of a cloud-first world:

- 1. The shift to cloud-native applications:** Applications are being rewritten and new ones developed—not for legacy environments where relatively static workloads are the norm, but for dynamic, scalable cloud environments. Motivating this shift is an implicit expectation that cloud and container infrastructure environments will deliver enhanced agility and improved ROI. This shift demands infrastructure transparency and real-time monitoring and analytics. Without these key pieces, neither applications nor their underlying plumbing can deliver the low-latency user experience end users have come to expect.
- 2. The shift to a DevOps philosophy:** Enterprises are moving away from IT silos and forming collaborative apps-driven cloud teams to speed time to market of enterprise software applications. This cultural transition has emphasized that application performance is a shared responsibility across both infrastructure operations and software development, and it demands a new set of tools that promotes collaboration, offers real-time data analysis, and enables self-service infrastructure management.

Enterprises embracing this new environment of DevOps and realizing the competitive advantage gained from software innovation are requesting three key attributes:

- 1. Efficiency.** Enterprises and telecoms want to streamline cloud infrastructure management to deliver high availability, high performance, and resource optimization.
- 2. Agility and Innovation.** Enterprises and telecoms want to provide application developers and operators with tools that reduce friction and enable rapid success, replacing traditional IT silos with DevOps-style collaboration.
- 3. Improved ROI.** Enterprises are looking for ways to operate large multicloud or hybrid cloud environments in the most cost-effective way, while delivering optimum performance.

Modern cloud management tools must offer enterprises and telecoms a "sweet spot" where monitoring, platform management, orchestration, security, and compliance intersect. This new category of management capabilities for the DevOps era can be referred to as service optimization.



The Juniper Networks-Intel Cloud and NFV Infrastructure Smart Orchestration Solution

In collaboration with Intel, AppFormix has integrated its cloud service optimization software with the Intel Resource Director Technology (RDT) hardware feature set. By using policy-based resource control, AppFormix removes the guesswork involved in allocating resources and allows operators to manage their infrastructure efficiently. For the first time, the AppFormix software in conjunction with the Cache Monitoring Technology (CMT), Memory Bandwidth Monitoring (MBM), and Cache Allocation Technology (CAT) features available in the new Intel Xeon E5 v4 processor family make it possible to enforce isolation between workloads and to detect and mitigate noisy neighbors in real time.

Cloud infrastructure is shared infrastructure, and therefore resource contention is inevitable. Whether an environment is built on a cloud platform such as OpenStack or a container runtime such as Docker, the physical infrastructure is divided across multiple applications and often across multiple tenants. AppFormix provides better isolation between workloads, better visibility into performance, and a collaboration space that keeps everyone on the same team. The result is a cloud environment that is predictable, reliable, and responsive to the needs of applications.

Through its unique "smart agent" approach to telemetry and analysis, AppFormix offers deep, meaningful insights in real time by analyzing the system's most accurate, high resolution, and relevant data directly at the source—all while consuming fewer resources than other monitoring systems.

The AppFormix software yields better performing applications as well as happier application developers and operators. With AppFormix, an operator can deliver a self-service IT experience to an application developer. App developers are able to consume infrastructure analytics within their applications and schedule workloads to achieve maximum performance. Operators manage

access with user permissions, granting developers the ability to see their infrastructure, set alarms, monitor, and troubleshoot application performance issues on their own, which in turn significantly reduces the load on the help desk.

Dashboards and APIs provide visibility into each layer of the stack—physical hardware, application software, and platform and management layers—to show users how applications consume server, storage, and networking resources in real time.

Real-time analysis enables the best possible decision making, especially when it comes to orchestrating resources and detecting faults in the infrastructure. The dynamic nature of cloud-native applications coupled with the shift to continuous deployment means that the demands placed by the applications on the infrastructure are constantly changing. The AppFormix automated, policy-driven resource orchestration and optimization features remove the guesswork involved in allocating resources and allow operators to manage more servers with efficiency. Proactive remediation results in a software-defined infrastructure (SDI) that is more responsive to the needs of the applications, yielding better performing applications as well as more productive application developers and operators.

Applications make demands on infrastructure, and the infrastructure makes promises to deliver. AppFormix is the software that makes this interaction work in the most demanding production environments.

Solution Components

Juniper Networks AppFormix Platform: AppFormix provides real-time control and monitoring of your cloud infrastructure without high-touch modifications to sensitive ecosystem layers, and especially without slowing down production services.

The system is designed around a distributed analysis engine, called the AppFormix Data Platform. This platform performs configurable, real-time evaluation of in-depth, high resolution metrics of the physical infrastructure and the virtualization layer




Figure 1. AppFormix solution architecture

executing VMs and containers. Distributed evaluation of metrics generates events on a message bus for prediction, optimization, and correlation analysis across the infrastructure. Users access the real-time metrics, events, and analysis results via a unified dashboard or by API.

“Using Intel’s game-changing technology, Resource Director Technology, we’re giving operators and developers unprecedented visibility into and control over how applications run at the processor level, solving thorny real-world problems like noisy neighbors in real time.”

Sumeet Singh, VP AppFormix Product and Founder, Juniper Networks

The primary subsystems are the Policy Controller, Data Platform, and Analytics. The Policy Controller manages policies for resource monitoring, analysis, and control. It also provides role-based access control. The Data Platform is a distributed system for metrics collection, analysis, and event generation. The Data Platform is fully cloud-aware and capable of associating raw resource usage data with elements of the infrastructure, virtualization, and application layers in a cloud environment. Analytics modules analyze metrics and events produced by the Data Platform. Analytics modules can further analyze data across multiple elements and provide higher level signals and information about the state of the SDI to operators and developers.

The AppFormix system integrates with platforms such as OpenStack, Docker, and Kubernetes via Adapter modules that discover physical and virtual elements in the environment. It configures those elements in the Policy Controller, and as an element enters and exits the system, the Adapter updates the Policy Controller.

Juniper Networks AppFormix User Interface: AppFormix exposes its functionality in two ways—REST-based APIs and a web-based Dashboard. Dashboard is a graphical web-based client that runs in a browser. Dashboard enables the user to configure policies on the SDI and offers users a visual representation of resource metrics, alarms, health, and reports. It displays the relationship between entities (hosts, instances, projects, aggregates), and allows navigation of data across different axes. For example, the consumption of an instance may be viewed in the context of the host on which it executes and competes for resources with other instances. Or, the instance may be viewed on a chart with other instances in the same project, even when the instances all execute on different hosts.

Juniper Networks AppFormix Policy Controller: AppFormix Policy Controller is a central policy engine that configures other AppFormix components. User configuration is stored as policies that determine how monitoring, analysis, and control apply to selected elements (compute jobs, containers, VMs, etc.) in the

software-defined data center (SDDC). As elements enter and exit the system, the Policy Controller configures the AppFormix components automatically according to current policies. For example, when a new virtual machine is created, the Policy Controller applies configuration for all policies for the tenant that owns that VM. The Policy Controller also provides role-based access control. Authentication is provided by integration with existing identity services, such as Keystone in OpenStack.

Juniper Networks AppFormix Data Platform: The AppFormix Data Platform handles all resource data in the system, from raw metrics to configurable alarms. Using a distributed architecture, the Data Platform collects, processes, stores, and publishes resource metrics at multiple time scales.

On each host, the AppFormix Agent collects and evaluates metrics relevant to a policy configured for elements on the host: the host itself, containers, virtual machines (VMs), and applications executing inside the VMs. From the raw metrics data, the Agent publishes meaningful events and telemetry that are associated with elements in the SDDC. By processing metric data at the site of collection, the Agent efficiently balances real-time evaluation at fine granularity with publishing at coarser time scales. Further, the distributed processing performed by multiple Agent instances scales with the size of the infrastructure.

Users can view real-time metrics in the Dashboard for current status and incident debugging. Summarized usage data is used to generate reports that assist users with capacity planning and chargeback. Alarms are generated in response to conditions defined in user-configured monitoring policies. Data from each category is streamed in real time on the message bus and stored in a database for retrieval.

Alarms are a configurable feature of the Data Platform that record when resource usage matches user-configured criteria. Alarms offer high-performance, fine-grained identification of conditions without shipping large quantities of raw metrics off the host (saving network bandwidth). All alarms are actionable and can be configured by the Policy Controller to effect change in the running state of the infrastructure; for example, dynamically changing how much cache is allocated to each container running on the host.

Juniper Networks AppFormix Analytics

AppFormix Analytics is a modular subsystem that provides higher order features to users. An analyzer module interfaces with the Data Platform and Policy Controller using REST-based APIs. An analyzer may listen for real-time data on the Data Platform’s message bus, access historic data in the database, or both. In addition, an analyzer may publish data back into the Data Platform. An analyzer may also set policy in the Policy Controller, to configure resource monitoring or to orchestrate resources via a control policy. AppFormix ships with several analyzers: Health, Reports, and Notifications.

- The Health analyzer produces a risk rating to help users identify elements that are not meeting user-defined SLAs.
- The Reports analyzer produces a report of resource consumption during a time period specified by the user. These reports are useful for capacity planning or chargeback. The reports are stored in the Data Platform and can be shared by users for collaboration.
- The Notifications analyzer posts Alarms to third-party notification services, or an incident management system. This configurable module can send Alarm data to any external REST-based service.

Intel Resource Director Technology (RDT): As introduced in the Xeon E5 v4 family of server processors, RDT provides both monitoring and control (allocation) capabilities. The monitoring features enable better characterization of the resource requirements of applications, for instance how much cache is needed to run at a certain level of performance, and can help detect “noisy neighbor” applications within the data center, which may slow down other higher priority workloads. In a complementary fashion, the allocation features help provide more control for prioritizing workloads, controlling noisy neighbors, and improving performance determinism in complex and dynamic environments with many active processor cores. The constituent hardware technologies described below are leveraged extensively by the AppFormix software suite to monitor and control the resource usage of containers and virtual machines.

Intel visibility into last-level cache (LLC) utilization per thread, app, container, or VM: Cache Monitoring Technology (CMT) provides new insight by monitoring the LLC utilization by individual threads, applications, or VMs. Example CMT usages include improved dynamic application characterization, “noisy neighbor” detection, performance debugging, advanced real-time provisioning, and resource-aware scheduling decisions.

Intel improved visibility into memory bandwidth utilization: Memory Bandwidth Monitoring (MBM) is an extension of CMT,

providing per-thread memory bandwidth monitoring for all threads simultaneously for the first time—enabling multiple VMs or apps to be independently tracked. New capabilities enabled include detecting “noisy neighbors” which overutilize memory bandwidth, characterizing and debugging the performance of bandwidth-sensitive applications, and more effective NUMA-aware scheduling, as visibility into both local and remote (e.g., QPI) bandwidth usage is provided.

Intel control over LLC usage and prioritization: Cache Allocation Technology (CAT) improves control by prioritizing important data center VMs, containers, and applications through software guided redistribution of LLC capacity. This ensures runtime determinism; for instance, protecting important VMs, virtual switches, or Intel Data Plane Development Kit (Intel DPDK) packet processing applications from resource contention, and preventing noisy neighbor interference across priority classes of workloads. Priority classes can also be created within the data center; for instance, prioritizing interactive workloads over batch compute jobs.

Intel improved isolation in code and data caching: Code and Data Prioritization (CDP) extends CAT to enable separate control over code and data placement in the last-level (L3) cache. Certain specialized types of workloads benefit with increased runtime determinism, enabling greater predictability in application performance.

How AppFormix on Intel Xeon E5 v4 Platform Works

Software-defined infrastructure presents new challenges with respect to resource sharing. Prior to Intel Resource Director Technology, resources such as L3 cache and memory bus bandwidth could not be monitored or controlled. With Juniper Networks AppFormix and the Intel RDT found in the Intel Xeon E5 v4 platform, it is possible to monitor and identify anomalous resource utilization behavior and take mitigating action.




Figure 2. Examples of real-time monitoring data as presented by the AppFormix Dashboard

The AppFormix Agent automatically detects if Intel RDT features are available on a host. For a host equipped with Intel Xeon E5 v4 processors, the AppFormix Agent is able to capture performance counters, such as L3 cache usage, cache misses, memory bandwidth, and instructions per cycle, provided by traditional performance monitoring counters as well as the new CMT and MBM monitoring technologies. The AppFormix Agent intelligently utilizes raw hardware counters and associates their values to elements of SDI such as OpenStack virtual machines, Docker containers, Kubernetes pods, and native host workloads.




Figure 3. AppFormix combined with Intel's RDT improves visibility and control over how applications run together on server platforms, allowing enforcement of priority levels, advanced provisioning, and mitigating "noisy neighbors" in a fully dynamic environment.

As an example, consider a noisy neighbor that is consuming a disproportionate share of the L3 cache. The AppFormix Agent can generate an alarm when this instance is detected. AppFormix Health analytics will receive the alarm on the message bus, and in response, may configure a policy that partitions the L3 cache to reduce usage by the noisy instance and thereby maintain the SLA for high priority applications. Alternately, the AppFormix Policy Controller enables policies to be proactively defined to manage resource utilization for all instances, such that high priority tenants have guaranteed resource allocations required to meet performance SLAs. The host-local, policy-aware AppFormix Agent can react in real time to ever changing conditions to ensure policy compliance.

The CAT feature of the Intel Resource Director Technology offers cache allocation control in the form of a) absolute MB cache usage, b) percentage of cache usage, or c) overlapped or isolated

cache regions. AppFormix exposes these controls in the policy for an instance, be it a VM or container. Thus, a policy may be created to provide consistent SLAs and performance guarantees to applications.

The AppFormix Health module uses a simple REST interface of the Policy Controller to dynamically update the cache allocation of an instance.

The following sample curl command depicts the API call:

```
curl -i \
  -H 'Content-Type: application/json' \
  -X PUT \
  -d '{ "CacheAllocation": 5 }' \
  http://<PolicyController>:7000/appformix/v1.0/instance/<instance-id>
```

Experimental Results

The performance benefits of the Intel CAT feature can be demonstrated using two back-to-back connected servers.

The load generator runs on Server-1, which is a previous generation platform including a two socket configuration based on Intel Xeon L5520 processors at 2.27 GHz. Each socket has 8 logical threads (SMT-enabled) and 8192 KB of L3 cache. The load generator requests the same file using 22 processes and 24 concurrent connections.

Server-2 is the system under test, which consists of an NGINX webserver with a total of 88 threads evenly distributed on the two socket Intel Xeon CPU E5-2699 v4 @ 2.20 GHz. Each socket has 44 logical threads (SMT on) and 56,320 KB of L3 cache. The NGINX processes serve files of different sizes: 10 MB, 1 MB, 100 KB, 10 KB.

The noisy entity used for the purposes of the experiment is a standard benchmark called Stream2,* and it runs 11 processes per socket evenly distributed on each Server-2 socket. The CAT masks are applied on both sockets via a cgroups kernel patch in an automated fashion using the AppFormix framework. The effective CAT mask applied to contain the "noisy neighbor" threads is 0x0003, mapping to approximately 10% of the cache, while the webserver threads use the full cache.

Performance results are shown below for two key metrics—throughput and effective webserver response time as observed at the load generator. Both metrics are improved. Absolute performance is improved by up to 27% by containing the "noisy neighbor" as shown in Figure 4.




Figure 4. Using the AppFormix software suite and Intel’s CAT to contain a “noisy neighbor” and prioritize the NGINX webserver yields up to a 27% performance increase.

Latency metrics also improve. For instance, the average response time of the server to a given request reduces up to 51% in the presence of a noisy neighbor as shown in Figure 5.




Figure 5. Using the AppFormix software suite and Intel’s CAT results in up to a 51% reduction in average response time (latency), improving the experience for end users.

Features and Benefits

Table 1. AppFormix and Intel Resource Director Technology (RDT) Solution Highlights

Solution Feature	Benefits
AppFormix Analytics	<ul style="list-style-type: none"> AppFormix software, in conjunction with the Intel Resource Director Technology (RDT), is available in the new Intel Xeon E5 v4 processor family, making it possible to detect noisy neighbors in real time. This is a strong enabler for collaboration between developers and operators.
AppFormix Control	<ul style="list-style-type: none"> AppFormix software, in conjunction with the Cache Allocation Technology (CAT) instruction set, makes it possible to enforce isolation between workloads and to detect and mitigate noisy neighbors in real time. This improves application performance, and better infrastructure improves ROI.

Summary—Juniper and Intel Deliver a Software-Defined Networking Stack

The combination of AppFormix orchestration and analytics software with the latest Intel Resource Director Technology (RDT) hardware features can yield substantial benefits in software-defined infrastructure (SDI). It allows prioritization of important containers, VMs, or applications; detection of and mitigation of noisy neighbors; and improving the workflow between infrastructure maintainers and software developers in a modern collaborative DevOps environment. Benefits in the general data center and communications applications are possible as well.

Next Steps

If you would like to learn more about this joint solution, please contact your Juniper Networks representative for more information.

About Intel

Intel makes possible the most amazing experiences of the future. You may know us for our processors. But we do so much more. Intel invents at the boundaries of technology to make amazing experiences possible for business and society, and for every person on Earth.

Harnessing the capability of the cloud, the ubiquity of the Internet of Things, the latest advances in memory and programmable solutions, and the promise of always-on 5G connectivity, Intel is disrupting industries and solving global challenges. Leading on policy, diversity, inclusion, education, and sustainability, we create value for our stockholders, customers, and society.

About Juniper Networks

Juniper Networks challenges the status quo with products, solutions and services that transform the economics of networking. Our team co-innovates with customers and partners to deliver automated, scalable and secure networks with agility, performance and value. Additional information can be found at [Juniper Networks](#) or connect with Juniper on [Twitter](#) and [Facebook](#).

Corporate and Sales Headquarters

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, CA 94089 USA
Phone: 888.JUNIPER (888.586.4737)
or +1.408.745.2000
Fax: +1.408.745.2100
www.juniper.net

APAC and EMEA Headquarters

Juniper Networks International B.V.
Boeing Avenue 240
1119 PZ Schiphol-Rijk
Amsterdam, The Netherlands
Phone: +31.0.207.125.700
Fax: +31.0.207.125.701



Copyright 2017 Juniper Networks, Inc. All rights reserved. Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners. Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

JUNIPER
NETWORKS