

Zenoss Service Dynamics 5.0

Architecture Overview

Table of Contents

INTRODUCTION	1
FLEXIBLE, ENTERPRISE SCALE PLATFORM	1
NEW SERVICE-ORIENTED PLATFORM IMPLEMENTATION	2
CENTRALIZED PERFORMANCE DATA STORAGE	3
SIMPLIFIED OPERATIONS WITH CONTROL CENTER.....	3
ZENOSS SERVICE DYNAMICS 5.0: SERVICES OVERVIEW	5
Collectors.....	5
Hubs.....	5
Apache HBase and Open TSDB.....	6
Unified Web Interface.....	6
Unified Monitoring	6
Event Management.....	7
Service Impact.....	7
Analytics	7
Scalability, High Availability, and Hardware Specifications	8
CONCLUSION	8

Introduction

The demand for agile, efficient business services is higher than ever, driving increased investment in virtualized, cloud, and converged infrastructure technologies. While incredibly dynamic, these technologies also bring with them a healthy amount of complexity. For IT operations teams, that adds several degrees of difficulty when trying to ensure service reliability.

The modern architecture in Zenoss Service Dynamics 5.0 was designed to accommodate the unique requirements of monitoring today's service infrastructures. As in previous releases, it helps reduce complexity by providing a unified, end-to-end view of all the interconnected resources and dependencies supporting service delivery – across physical, virtual, or cloud-based environments. The 5.0 release, however, introduces a new service-oriented architecture that makes your monitoring just as dynamic as your real-time infrastructure. By abstracting services so that they can share compute and storage resources, you can drive more rapid implementation of monitoring, streamline updates and upgrades, build out high availability, and scale up or out more easily.

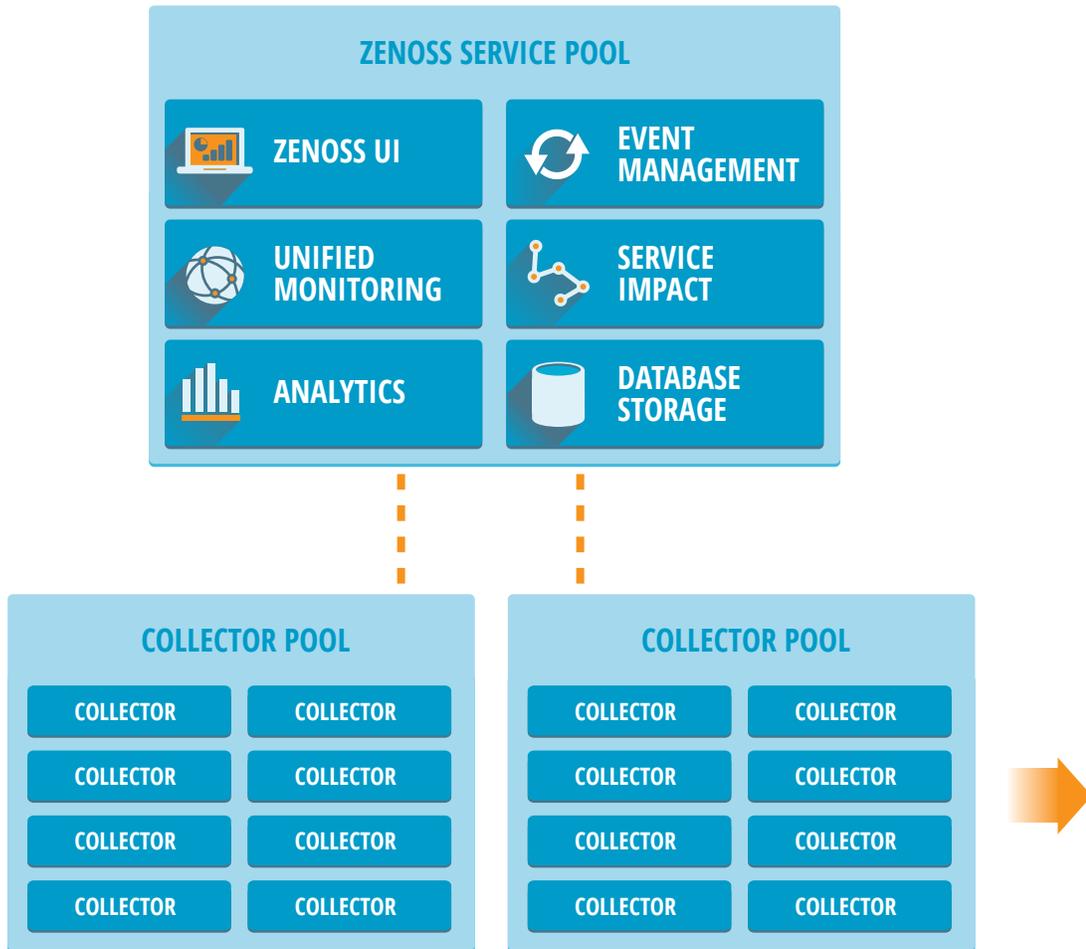
Flexible, Enterprise Scale Platform

Zenoss Service Dynamics has always been designed from the ground up as a flexible, highly scalable, and extensible architecture for addressing end-to-end service assurance challenges. With the release of Zenoss Service Dynamics 5.0, customers get more flexibility, scalability, and extensibility than ever before. All of the improvements in 5.0 build on core Zenoss attributes that enable monitoring across multiple platforms in dynamic environments, whether physical components (servers, switches, routers, SANs), virtual components (VMs, VLANs), or private and public cloud platforms (Cisco UCS, VCE Vblock, NetApp FlexPod, AWS, OpenStack). These foundational attributes include:

1. **Unified and Agentless:** Zenoss Service Dynamics offers a unified platform for comprehensive monitoring of hundreds of devices, systems, and applications in dynamic, highly distributed environments. In these conditions, agent approaches break down quickly, as the time necessary for deploying, updating, and integrating agents undermines the ability to accurately and efficiently monitor service reliability. By using agentless collection and control techniques — such as secure access methods, management APIs, and synthetic transactions — Zenoss Service Dynamics is able to provide the real-time, end-to-end visibility necessary to keep services up and performing.
2. **Cloud-era Scalable:** In dynamic environments, you need to scale up and scale down elastically to meet demand. Zenoss Service Dynamics can scale to manage hundreds of thousands of devices in globally distributed configurations.
3. **Open and Extensible:** Zenoss Service Dynamics allows you to rapidly extend or customize monitoring for any resource in your environment, as well as integrate with other management tools. Organizations can use Zenoss-provided ZenPacks, create their own ZenPack plugins, or leverage any of the publicly available ZenPacks from the global Zenoss community of more than 100,000 developers and partners.

New Service-oriented Platform Implementation

Previously, the Zenoss platform was deployed in a standard server application model, with one server for unified monitoring and event management, a separate server for service impact management, and a third server for analytics. Agentless collectors were deployed on one or multiple servers depending on how many resources were being monitored. In the 5.0 release, Zenoss has moved to a dynamic, service-oriented model where all capabilities are deployed and accessed in resource pools (see Services Overview for descriptions of main service components).



ONE EXAMPLE OF HOW ZENOSS SERVICE DYNAMICS 5.0 RESOURCE POOLS CAN BE CONFIGURED. RESOURCE POOLS CAN BE GROUPED IN ANY CONFIGURATION TO FIT YOUR ORGANIZATIONAL REQUIREMENTS.

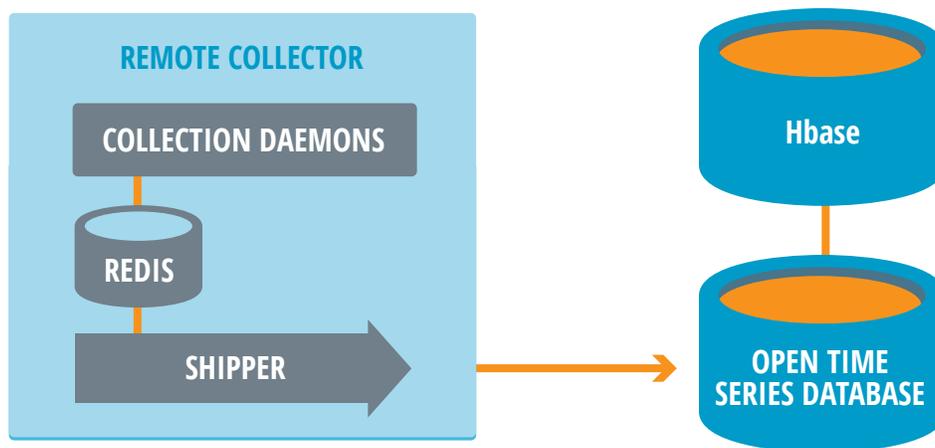
Resource pools dynamically share compute and storage resources, making it easier for Zenoss Service Dynamics to scale up and down to meet monitoring demands. Services are assigned to a pool and are dynamically distributed across all of the hosts assigned to that pool. If a service is overloaded, you can add new hosts to the resource pool and increase the instances of a service, which will be distributed across the new hosts.

Centralized Performance Data Storage

Another key difference in the 5.0 architecture is the ability to centralize performance data storage. Previously, performance monitoring data was stored on Round Robin Database (RRD) files on each collector. That required more expensive, high performance storage hardware for remote collectors. It also meant that you could only store a limited window of raw, real-time performance data – because as soon as the RRD files filled up, it started overwriting the oldest data. If a remote collector failed, you would also lose all of the performance data in its RRD – to implement failover, you'd need to deploy duplicate servers.

With a new Big Data architecture, Zenoss Service Dynamics 5.0 retires the RRD storage model. Instead, all collectors (still agentless) leverage centralized storage, which provides both performance and reporting benefits. Real-time data is cached locally to a Redis, then continually streamed to an Open Time Series Database (Open TSDB). The Open TSDB then writes to Apache HBase, giving you the flexibility and power of a Big Data solution for reporting at the most granular levels in any window of time you choose. Unlike the previous RRD model, no data gets aggregated or written over without your express consent.

Finally, RRD data was previously limited to a set interval, with a default of 5 minutes. With the new 5.0 architecture, collection intervals can be changed at any time. This means that if a threshold or issue triggers an event, the data collection interval can be automatically adjusted to collect more often so you can capture more information about the event. Once the event is cleared, you can reset the data collection interval to the default.

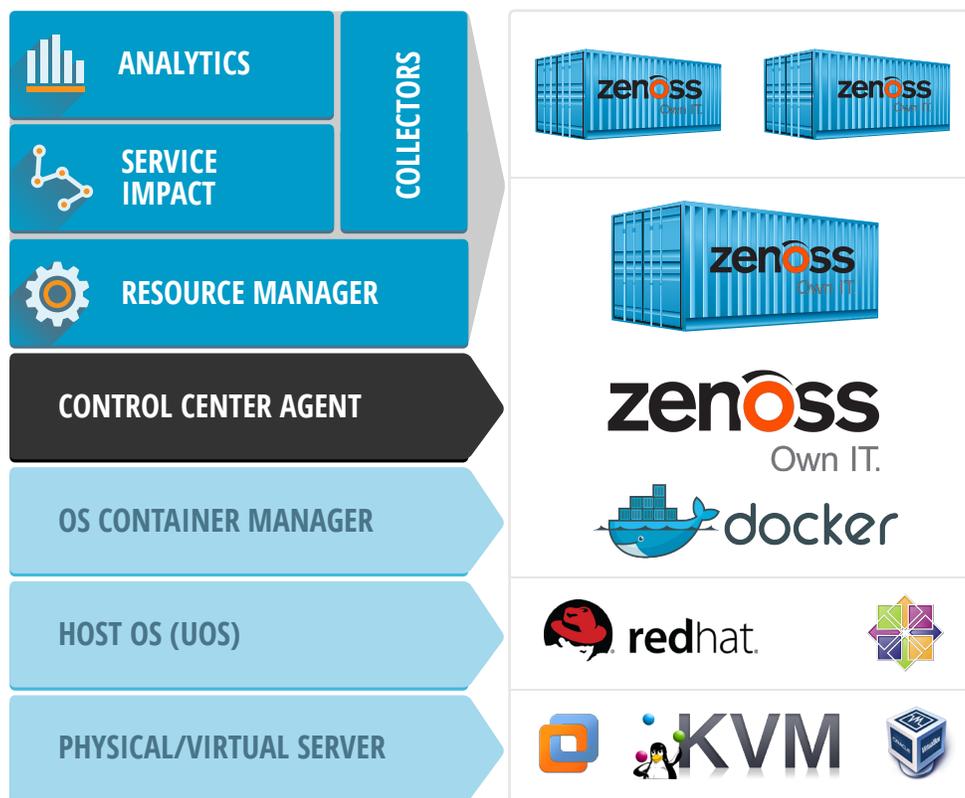


Simplified Operations with Control Center

In addition to being more scalable from a technology standpoint, Zenoss Service Dynamics 5.0 operations becomes simpler and more scalable by introducing Control Center. Control Center is an application service orchestrator that provides out-of-band management services. Built on Docker, it allows you to run Zenoss components inside containers, packaging the application with any and all dependencies, including prerequisites like the appropriate versions of RabbitMQ, MariaDB, and Python. This removes any possibility of software conflicts within a package since each container is isolated from everything else running on that machine. Containers allow Zenoss to be deployed more efficiently and create a consistent running environment. This consistency reduces operational overhead, which is crucial when attempting to monitor a cloud-scale service infrastructure.

Under the covers, Control Center uses a union file system, which allows you to take snapshots, roll back, and push out changes in a much more agile fashion. The union file system consists of a set of layers, making it easier to coordinate updates and upgrades. Each update builds on the previous layers of code, meaning when you apply a patch or update, you are only modifying one layer. Once a change is committed, the Control Center will push that change as a new layer out to all of your hosts, which look for and install only that layer. If any issues occur during installation, you simply uninstall that layer to roll back the application to its previous state.

Finally, Control Center also provides a platform for efficiently administering your Zenoss Service Dynamics 5.0 implementation. It allows you to assign and relocate services as needed, spin up new collectors with a few clicks, and track metrics for hosts and services. While not an alerting tool, if you do experience a failure or degradation in your Zenoss services, Control Center includes central logging for all Zenoss services and can help you determine the source of the issue and get monitoring back on line as quickly as possible.



ANATOMY OF THE CONTROL CENTER

Zenoss Service Dynamics 5.0: Services Overview

The Zenoss Service Dynamics product suite consists of two products: Resource Manager and Service Impact and Analytics. Resource Manager is the foundation, providing unified monitoring, inventory and modeling, and event management. Service Impact and Analytics adds the ability to identify how services are impacted by a performance degradation or fault in your infrastructure, as well as provide a central view into utilization trends, capacity planning, and operational anomalies that might impact service performance and availability.

Together, these two products deliver the services you need to provide unified service insight into your end-to-end service infrastructure, helping to improve service quality and reduce operational costs. This unified platform is possible because the Zenoss Service Dynamics products share a common services architecture, allowing them to work seamlessly to centralize control over any resource in your environment. An overview of these services is provided below.

Collectors

Collectors are an agentless way of gathering performance and availability data. Data collection is accomplished using a variety of protocols, such as WinRM, WMI, JMX, SNMP, SMTP, SSH, HTTP, Syslog and SQL, and can be customized to leverage data from proprietary or custom APIs. Having the ability to simultaneously collect data using all of these protocols makes collectors extremely efficient, spanning physical, virtual, and cloud-based infrastructure devices, components, or objects. Collectors also collect event data either directly from a device or through an element manager like VMware vSphere or NetApp ONTAP.

Each collector can scale up to approximately 100,000 data points for a standard five-minute polling cycle. You can also spin up multiple collector workers, expanding monitoring to accommodate one million data points every five minutes. Zenoss Service Dynamics 5.0 makes it possible to load balance devices across collector workers automatically, making it simpler and more efficient to scale. The number of collectors you will need to deploy depends on the size of your network and how often you schedule modelling updates and performance data collection. As mentioned above, if a collector is overburdened, adding new collectors and spinning up new collector workers takes only a few minutes.

Hubs

Hubs push device templates to collectors to define what information is collected from each managed resource. They also serve as the pass-through point for all event data. Status and availability information, such as ping failures and threshold breaches, are returned through the hub to the event processing system, where transforms and triggers will kick off notifications if appropriate. Depending on your environment, you can have a single or multiple hubs, which would typically be deployed in the primary resource pool. This flexibility allows you to scale as needed to fit your environment.

Apache HBase and Open TSDB

Zenoss Service Dynamics 5.0 allows customers to take advantage of Big Data performance, availability, and scale by implementing Cloudera Enterprise (with Apache HBase at its core) as a backend data store. With this new Big Data backend, it is possible to store a virtually unlimited amount of performance and availability data at the resolution it was collected at – no rollup required. This can give clear, granular insight into historical trends, as well as provide long-term storage.

On the front end, Zenoss Service Dynamics 5.0 uses Open TSDB for reading and writing at scale to the Big Data backend. Unlike traditional relational databases, the Open TSDB can easily handle the high transaction volume of time series data required in a real-time monitoring environment. In addition to handling high volumes of complex data, the Open TSDB also provides a performance processing pipeline that allows you to scale the solution effectively.

Unified Web Interface

The Zenoss Service Dynamics web interface is a unified user interface that sources data from any Zenoss service. It serves as a single pane of glass to provide real-time views into all of your environment's resources, centrally managing events, determining relationships between services and infrastructure issues, and analyzing operational performance.

The unified interface provides operational transparency, giving administrators access to a single, consistent information set about the environment, regardless of their organizational role. Each administrator can customize the interface using portlets, making it easy to view the information most important to their role. The interface is a collection of services, so it can easily scale to accommodate hundreds of concurrent users.

Unified Monitoring

Unified Monitoring provides the underlying inventory, modeling, and monitoring of resources that make it possible to view the reliability of your end-to-end service delivery infrastructure in the unified web interface. Unified Monitoring discovers devices and populates a resource model – which serves as a complete inventory of your servers, storage, network devices, and applications down to the component level (interfaces, services, and processes). Managed resources can be added by a discovery process or via APIs.

Unlike traditional CMDBs that rely on batch process updates, the Zenoss resource model is maintained in near real-time so you have an immediate understanding of how each device or component is working in your environment. Some individual device components, like VMware VMs or VLANs on a Nexus switch, can automatically update the resource model via configuration change events. Other device components, such as physical interfaces on a Cisco 3600-series switch or Linux server file system, are added and removed during regularly scheduled modeling processes.

Once a resource is added to the model, Zenoss begins monitoring it immediately. Monitoring templates are assigned and distributed to the appropriate collectors via a hub.

Event Management

Zenoss Service Dynamics also provides robust event management for all managed resources. The event engine is capable of processing an extremely high volume of data – with one event processor able to handle more than 100 million events daily.

To help administrators effectively manage events, the system aggregates, filters, de-duplicates, and masks events to ensure that only those events that pose a risk to service delivery are surfaced. During event storms, the event system is aided by root-cause analysis and confidence ranking engines in Service Impact, which help prioritize events that pose the highest risk to service delivery. The additional information provided by these service events enables the event system to quickly parse through thousands of console events in just seconds.

Service Impact

Service Impact houses the service model where relationships and dependencies are mapped between managed resources and the services they support. Service constructs are defined based on logical business constructs to identify the infrastructure groupings that support a specific application service. For example, a CRM application service might require email, web, and database IT services to operate. For the CRM service to be considered functional, this collection of supporting services has to be factored into that logical construct. These constructs are then used for calculating the most likely root cause of service issues, evaluating the impact of service downtime, and aiding in event aggregation, filtering, de-duplication, and prioritization.

Once the logical service hierarchy is defined, Service Impact uses advanced modeling capabilities to pull in all relevant infrastructure elements – already defined as part of the Unified Monitoring service. These resources – whether VM partitions, blades, chassis, storage or networking interfaces – are discovered and mapped into the relevant service dependency graphs. It uses this information to process events associated with these resources against root-cause analysis and confidence ranking engines to generate service events that pinpoint the most likely source of a given incident. These events help alert administrators to service risk, decrease triage time, and expedite root cause analysis.

Analytics

Analytics directly leverages the resource and service models to report on service levels and risk assessment. The open business intelligence engine that delivers this functionality creates scalable and rich analytics for both historical reporting and trending information. Analytics can be deployed as part of the primary Zenoss Service Pool, or on its own pool, giving organizations the flexibility to make use of resource intensive reporting and analytics features in whatever way introduces the least impact on monitoring operations.

The Analytics architecture combines extract, transform and load (ETL) procedures with a data warehouse and reporting tools to process and report data. Analytics aggregates data over time and is meant for long term trending. It will purge data from the data warehouse at regular intervals, rolling up hourly or daily aggregates on whatever schedule you define.

This aggregate information is critical to understanding health and utilization trends, allowing you to forecast capacity needs and anticipate availability and performance problems. You can draw conclusions about upcoming operational issues and infrastructure restraints that could impact service delivery by examining service availability, service performance, or CPU usage over time and using trend extrapolation.

Scalability, High Availability, and Hardware Specifications

Zenoss Service Dynamics 5.0 is designed to be incredibly flexible, scalable, and configurable, meeting the specific needs of even the largest, most dynamic environments. Collectors scale easier than ever before, since you can now add workers for each collector with a few clicks in the Control Center. The new workers will then be distributed across the hosts in the pool. If storage or compute resources run low, you simply add another server to the pool.

High availability is also much easier to achieve, since when a single service fails, another automatically steps into take its place. There is no longer the need to have dedicated backup hardware for each service – and the Control Center can inform you what Zenoss services are up or down at any given moment.

In terms of the hardware requirements for deploying your highly scalable and available monitoring system, the flexibility of the Zenoss platform means that there is no one set of deployment guidance. Hardware requirements will depend on your specific environment, since any service can be deployed on either a physical or virtual machine in any number of resource pools. That flexibility doesn't mean complexity. Because services run in containers that isolate them from any other services running on that system, there is less conflict and much more consistency in how the service operates. To ensure that you have a detailed estimate of hardware requirements relevant to your specific infrastructure, the Zenoss professional service team conducts a personalized technical review for all new deployments.

Conclusion

Flexibility and scalability are a must for organizations wanting to adapt to managing real-time, dynamic infrastructures. Zenoss has always excelled at providing an enterprise-scale, cloud-era IT operations platform that maximizes infrastructure performance and service availability – and with Zenoss Service Dynamics 5.0 those benefits multiply. While still maintaining its open, agentless, and extensible properties, the new service-oriented architecture makes Zenoss Service Dynamics even more flexible, scalable, and efficient than ever before. Unified monitoring, event management, service impact, and analytics can be implemented more quickly, data can be stored and mined more broadly, and administrators can manage monitoring operations more effectively. All of these new capabilities combine to make it even easier to increase service quality while reducing operational costs.



11305 Four Points Drive
Building 1, Suite 300
Austin, TX 78726, USA

Phone: +1-512-687-6854
Toll-free: +1-888-936-6770
www.zenoss.com