

Virtana Compute Observability for Host Operating Systems

Mission-critical application-focused monitoring for physical, virtual and cloud servers

Virtana Compute Observability for Operating Systems is part of the Virtana full-stack hybrid infrastructure observability platform.

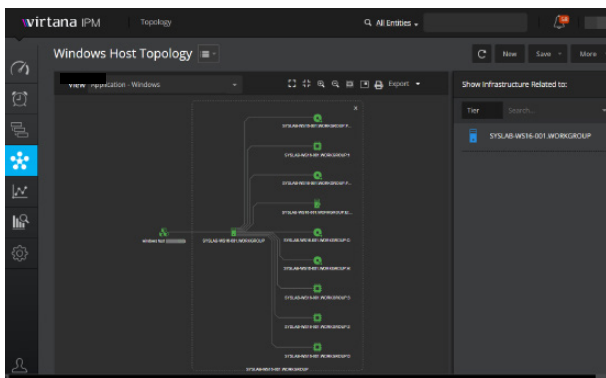


Figure 1: View of compute assets delivering application services.

Virtana Observability for Operating Systems

- **Prevent problems from occurring:** With AI and trend-based predictive capacity management within the full context of application topology and infrastructure.
- **Reduce MTTR:** With AI-powered runbook style investigations, automated service desk integration, and prescriptive best-practice resolution recommendations.
- **Get up and running quickly:** With application context-aware automated server discovery, automatically applied best-practice monitoring, predictive capacity management and dashboards

Critical challenges with monitoring and managing servers for mission-critical applications across hybrid infrastructure

Managing and monitoring servers in today's hybrid mission-critical application environments requires context, automation and analytics that go well beyond the simple metrics and trending used in the past, or available from cloud or virtualization providers.

Application and infrastructure visibility across entire hybrid infrastructure stacks

For any hybrid application, supporting servers need to be visible regardless of where they reside within the entire infrastructure set. What's more, visibility needs to include application infrastructure context. Not just the server itself, but it's context within the application's supporting topology, storage and network resources (cloud and local), virtual hosts and clusters.



Key capabilities of Virtana Compute Observability for Operating Systems

Automated discovery, depending mapping, monitoring and dashboarding

When even physical hosts can be repurposed as needed, and when new virtual and cloud servers come and go on demand or via automated orchestration the traditional methods of building and applying monitoring aren't an option. Server infrastructure needs to be continually discovered, while monitoring, alerting, dashboards and full application topology and views are automatically configured. And, this needs to happen across all locations where infrastructure lives, across; physical machines, virtual environments, storage, and cloud infrastructure.

AI-driven analytics that drive uptime – preventing problems before they happen and reducing MTTR when problems occur

Preventing problems and then resolving issues quickly requires tools that can analyze issues and usage, recognize trends and patterns, as well as use these tools to triage problems quickly while also balancing workloads for long-term usage. AI, machine learning, heuristics, trending and statistical methods applied across to servers, underlying infrastructure and full application stacks are needed.

Automated server discovery and visibility for hybrid applications

- Continuous, automatic discovery of servers across your full infrastructure stack within the context of applications and environments. Sources for discovery of physical, virtual and cloud servers include AWS, Azure, virtual environments, NetFlow

data and enterprise storage connections. Application context extends this information via connections to AppDynamics, DynaTrace, and ServiceNow (additional CMDB support available via professional services).

- Automatically apply best-practice or customized monitoring, thresholds, and alerting for both status and capacity management of server's CPU, network, and storage.
- Automatically create best-practice dashboards and reports for discovered servers within the context of the full-application stack being monitored and managed.
- Server dashboards and reports automatically link to application specific topologies and displays.

Problem resolution

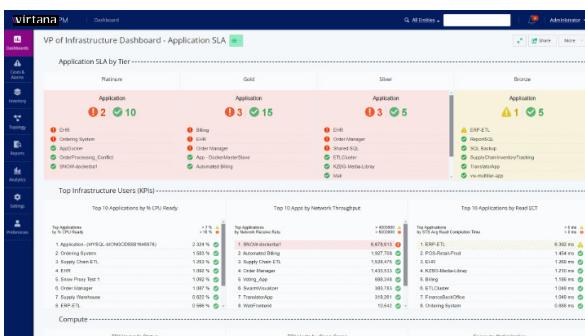
- The most granular, long-term datasets available combined with runbook-style automated investigations proactively identify problems with servers, the underlying infrastructure, and broader application environments that other solutions can't even identify.
- AI-powered recommendations that integrate easily with ITSM solutions such as ServiceNow enable quick resolution of the problem once identified – And include tools that enable automation of problem resolution.

Capacity management

- The longest term, most granular data sets in the industry for servers, underlying infrastructures and broader applications.
- Using this high-quality data, multiple AI, statistical, and trend-based analytics are applied to capacity forecasting to ensure server, infrastructure, and application availability.

Workload automation

- Real-time, AI-driven workload optimization recommendations for servers within the context of underlying infrastructure and applications.
- Workload optimization AI and statistical models capture and extend years of real-world application uptime and SLA delivery experience from Virtana.



Virtana Compute Observability for Operating Systems

KEY FEATURES

		Discovery & Mapping	Agentless Monitoring	Performance Metrics	Capacity Metrics	Best Practice Alarming	Custom Alarming	Intelligent Problem Resolution
Operating System onpremise	On-prem Windows Linux Solaris	✓	Compute, Memory, Network, Storage	✓	✓	CPU, Memory, Disk, Latency	✓	Guided Investigations
	AWS Windows & Linux	✓	Compute, Memory, Network, Storage	✓	✓	CPU, Memory, Disk, Latency	✓	Guided Investigations
	Azure Windows & Linux	✓	Compute, Memory, Network, Storage	✓	✓	CPU, Memory, Disk, Latency	✓	Guided Investigations
	Kubernetes	✓	Compute, Memory, Network, Storage	✓	✓	CPU, Memory, Disk, Latency	✓	Guided Investigations

		AIOps Driven Analytics			
		Application Service Assurance	Predictive Capacity Management	Workload Infrastructure Balancing	Problem Resolution & Avoidance
Operating Sytem	Windows Linux Solaris	✓	✓		✓
Virtualization	vSphere HyperV KVM PowerVM	✓	✓	✓	✓