



**Hewlett Packard**  
Enterprise

# **HPE Operations Manager i**

Software Version: 10.10

## **OMi Concepts Guide**

Document Release Date: 18 December 2015

Software Release Date: December 2015

## Legal Notices

### Warranty

The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HPE shall not be liable for technical or editorial errors or omissions contained herein.

The information contained herein is subject to change without notice.

### Restricted Rights Legend

Confidential computer software. Valid license from HPE required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

### Copyright Notice

© Copyright 2015 Hewlett Packard Enterprise Development LP

### Trademark Notices

Adobe® and Acrobat® are trademarks of Adobe Systems Incorporated.

AMD, the AMD Arrow symbol and ATI are trademarks of Advanced Micro Devices, Inc.

Citrix® and XenDesktop® are registered trademarks of Citrix Systems, Inc. and/or one more of its subsidiaries, and may be registered in the United States Patent and Trademark Office and in other countries.

Google™ and Google Maps™ are trademarks of Google Inc.

Intel®, Itanium®, Pentium®, and Intel® Xeon® are trademarks of Intel Corporation in the U.S. and other countries.

iPad® and iPhone® are trademarks of Apple Inc.

Java is a registered trademark of Oracle and/or its affiliates.

Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.

Microsoft®, Windows®, Lync®, Windows NT®, Windows® XP, Windows Vista® and Windows Server® are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

NVIDIA® is a trademark and/or registered trademark of NVIDIA Corporation in the U.S. and other countries.

Oracle is a registered trademark of Oracle Corporation and/or its affiliates.

Red Hat® is a registered trademark of Red Hat, Inc. in the United States and other countries.

SAP® is the trademark or registered trademark of SAP SE in Germany and in several other countries.

UNIX® is a registered trademark of The Open Group.

## Documentation Updates

The title page of this document contains the following identifying information:

- Software Version number, which indicates the software version.
- Document Release Date, which changes each time the document is updated.
- Software Release Date, which indicates the release date of this version of the software.

To check for recent updates or to verify that you are using the most recent edition of a document, go to: [https://softwaresupport.hp.com/group/softwaresupport/search-result?keyword=.](https://softwaresupport.hp.com/group/softwaresupport/search-result?keyword=)

This site requires an HP Passport account. If you do not have one, click the **Create an account** button on the HP Passport Sign in page.

## Support

Visit the HPE Software Support website at: <https://softwaresupport.hp.com>

This website provides contact information and details about the products, services, and support that HPE Software offers.

HPE Software Support provides customer self-solve capabilities. It provides a fast and efficient way to access interactive technical support tools needed to manage your business. As a valued support customer, you can benefit by using the support website to:

- Search for knowledge documents of interest
- Submit and track support cases and enhancement requests
- Download software patches
- Manage support contracts
- Look up HPE support contacts
- Review information about available services
- Enter into discussions with other software customers
- Research and register for software training

Most of the support areas require that you register as an HP Passport user and sign in. Many also require a support contract. To register for an HP Passport ID, go to <https://softwaresupport.hp.com> and click **Register**.

To find more information about access levels, go to: <https://softwaresupport.hp.com/web/softwaresupport/access-levels>

### HPE Software Solutions & Integrations and Best Practices

Visit HPE Software Solutions Now at <https://softwaresupport.hp.com/group/softwaresupport/search-result-/facetsearch/document/KM01702710> to explore how the products in the HPE Software catalog work together, exchange information, and solve business needs.

Visit the Cross Portfolio Best Practices Library at <https://hpln.hp.com/group/best-practices-hpsw> to access a wide variety of best practice documents and materials.

# Contents

Chapter 1: Welcome to this Guide .....	7
How this Guide Is Organized .....	7
Who Should Read this Guide .....	8
Chapter 2: Introduction to Operations Manager i .....	9
Operations Bridge for a BSM Solution .....	9
Consolidated Event and Performance Management .....	11
Correlating Events .....	13
Topology-Based Event Correlation .....	15
Structured Problem Solving .....	16
Management Packs .....	17
Scalable Architecture with Multiple Servers .....	19
Monitoring Automation .....	21
User Engagement .....	24
Integration Interfaces .....	24
Business Value Dashboard .....	25
User Roles and Responsibilities .....	26
Chapter 3: Operator Workflow .....	28
The Operator Environment .....	28
Other Roles .....	32
Chapter 4: Monitoring Developer Workflow .....	33
Initial Analysis .....	33
Define Health Indicators .....	33
Configure Monitoring Automation .....	34
Other Tasks .....	34
Other Roles .....	35
Chapter 5: IT Operations System Administrator Workflow .....	36
Installation and Configuration Tasks .....	36
Oversee the OMi Installation .....	37

Tune Infrastructure Settings .....	37
Configure Users and User Roles .....	37
Other Responsibilities .....	38
Ongoing Tasks .....	38
Operations Bridge .....	38
Other Roles .....	39
Chapter 6: Application Expert Workflow .....	40
Installation and Configuration Tasks .....	40
Ongoing Tasks .....	40
Other Roles .....	41
Summary .....	42
Send Documentation Feedback .....	43



# Chapter 1: Welcome to this Guide

This guide is an introduction to Operations Manager i, and how it enables you to improve the efficiency of your IT services and infrastructure.

## How this Guide Is Organized

This guide contains the following information:

- ["Introduction to Operations Manager i" on page 9:](#)  
A high-level overview of the most important features helps you understand how you can use Operations Manager i to improve the performance, availability, and efficiency of your IT environment.
- ["Operator Workflow" on page 28:](#)  
A description of a typical day for Dave, the IT Operations operator, and how he uses event management to prioritize his daily tasks.
- ["Monitoring Developer Workflow" on page 33:](#)  
A description of the role of Mike, an IT Operations monitoring developer, and how he monitors a new application.
- ["IT Operations System Administrator Workflow" on page 36:](#)  
A description of the role of Matthew, and how he oversees the Operations Manager i environment and configures the operational infrastructure to integrate all the applications and servers in his domain.
- ["Application Expert Workflow" on page 40:](#)  
A description of the role of Alice, and how she configures generic monitoring solutions for all the applications and servers in her domain.

## Who Should Read this Guide

You should read this guide if you are one of these users:

- An IT Operations operator
- A DB, Exchange, SAP, or other subject matter expert who designs the monitoring scenarios for these enterprise applications
- An IT Operations monitoring developer
- An IT Operations system administrator
- An IT Operations application administrator



# Chapter 2: Introduction to Operations Manager i

Read this chapter for a high-level overview of Operations Manager i, and how it enables you to improve the efficiency of your IT services and infrastructure.

This chapter includes an architectural overview, shows how Operations Manager i fits into an Business Service Management (BSM) solution, and describes the underlying concepts.

This chapter is structured as follows:

- ["Operations Bridge for a BSM Solution" below](#)
- ["Consolidated Event and Performance Management" on page 11](#)
- ["Structured Problem Solving" on page 16](#)
- ["Management Packs" on page 17](#)
- ["Scalable Architecture with Multiple Servers" on page 19](#)
- ["Monitoring Automation" on page 21](#)
- ["Integration Interfaces" on page 24](#)
- ["Business Value Dashboard" on page 25](#)
- ["User Roles and Responsibilities" on page 26](#)

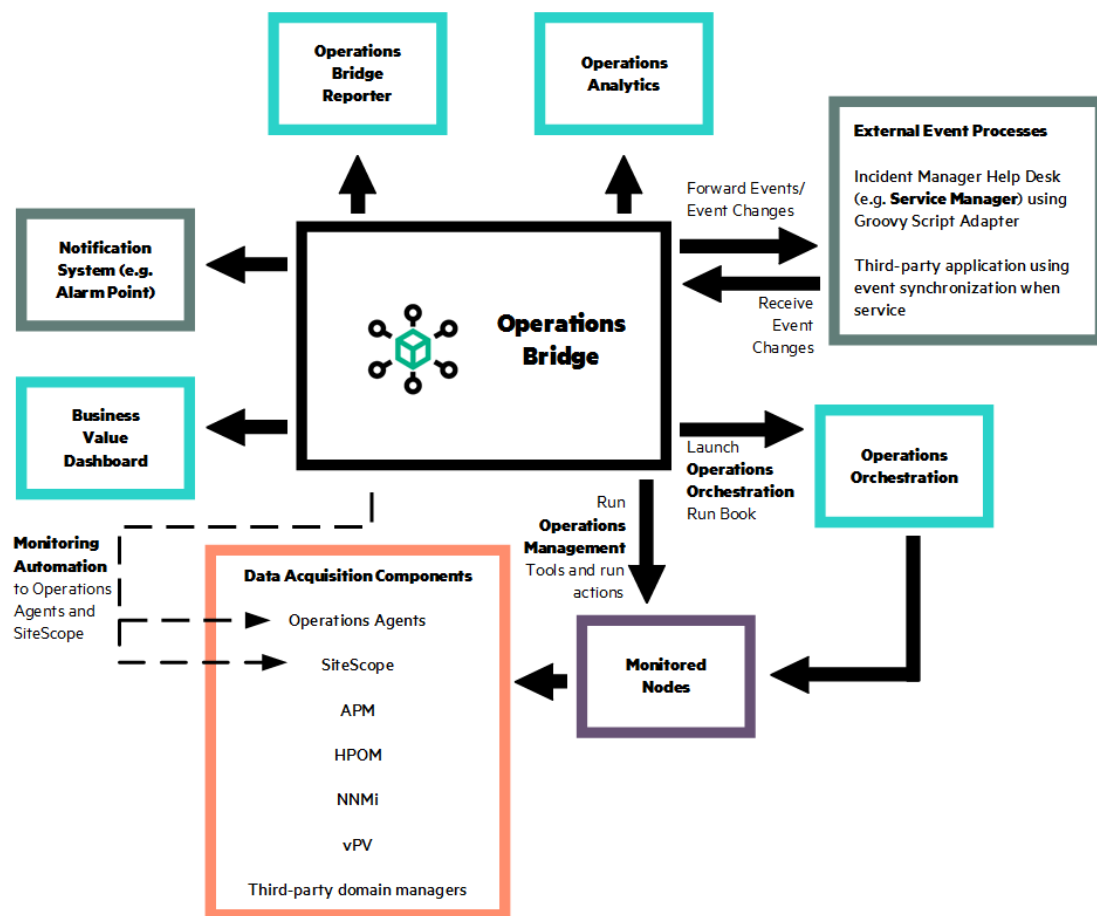
## Operations Bridge for a BSM Solution

OMi is the event management foundation for a complete Business Service Management (BSM) monitoring solution. As the operations bridge, it consolidates all IT infrastructure monitoring in a central event console, and relates the events to the IT services that depend on that infrastructure. Users benefit from a common structured event management model that applies the same processes to both business service management and IT infrastructure management.

OMi links infrastructure management with application and business service management. It combines events from HPE Application Management components, such as Business Process Monitor (BPM), Real User Monitor (RUM), and Service Level Management (SLM), with the events from the System and Network products, such as OM and HPE Network Node Manager i (NNMi). This enables you to keep track of all the events that occur in your monitored environment.

The figure below shows a typical deployment example where OMi is the operations bridge in a BSM solution. OMi provides automated monitoring and integration of multiple external applications.

All event and performance management originating from servers, networks, applications, storage, and other IT silos in your infrastructure are consolidated into a single event stream in an advanced, central event console. The console displays monitoring alerts to the appropriate team of operators.



All event and performance management originating from servers, networks, applications, storage, and other IT silos in your infrastructure are consolidated into a single event stream in an advanced, central event console. The console displays monitoring alerts to the appropriate team of operators.

You can quickly identify, monitor, troubleshoot, report on, and resolve problems in your distributed IT environment. These abilities make it possible for you to improve the performance and availability of the infrastructure and services in your monitored environment, adding to the efficiency and productivity of your business. OMi helps you locate and solve event-related issues before business service quality degrades. It offers the tools that help operators solve problems without involving a subject matter expert. This frees subject matter experts to focus on strategic activities.

### Data Acquisition from Multiple Sources

The events, regardless of where they originate, are processed and managed in a unified manner.

Examples of event sources include:

- HPE Operations Agents configured by OMi
- HP Operations Manager for UNIX with an HP Operations management server running on an HP-UX, SPARC Solaris, or x64 RHEL platform
- HP Operations Manager for Windows
- HPE Network Node Manager i (NNMi)
- HPE Business Process Monitor (BPM)

- HPE Real User Monitor (RUM)
- HPE SiteScope
- HPE Systems Insight Manager
- Third-party management software, normally used to monitor specific environments or special needs not monitored by other solution components, such as Microsoft Systems Center Operations Manager or Oracle Enterprise Manager. Connectors to integrate third-party management software, such as Microsoft SCOM, Nagios, and IBM Tivoli into HPE OMi are available from the [HPE Live Network Content Marketplace](#).

## Consolidated Event and Performance Management

The operations bridge is where events of all types from multiple sources are consolidated into a centralized console. “Perspectives” provide operators with different levels of information about the events they are responsible for. For example, general event handling is done in the Event Perspective, while the Health Perspective provides additional service health-related information about the events. These perspectives are centered around the Event Browser.

### Event Information

Events report important occurrences in the managed IT environment. They are generated by domain managers, forwarded to OMi, and then mapped to related configuration items (CIs) in the RTSM. These events are assigned to operators for resolution. In the Event Browser, operators can see a complete overview of all active events that need to be worked on. They can see such things as the event severity, the type and category of event, the source of the event, the time and location of the event, and the affected configuration item.

Events pass through a “lifecycle,” which is an informative way to display and monitor the status of an event. An operator’s workflow is based around the lifecycle of an event. The lifecycle state of an event represents the progress of the investigation into the problem that caused the event. An operator assigned to an event opens an investigation and works on finding a solution to the event’s underlying problem. Experts can then assess the proposed solution, verify that it solves the problem that caused the event, and close the event, which completes the lifecycle.

Operators can configure the Event Browser to suit the requirements of their typical workflows. The contents of the Event Browser are filtered according to the selected view or configuration item. Operators can configure new filters or modify existing filters, according to their needs, to change the information displayed. Filtering the Event Browser content helps operators focus on the most useful information, for example, to identify the highest priority events and to determine which of these events should be worked on first to minimize their impact on business services. You can also configure users and groups so that they can see only the events filtered by views associated with that user or group.

You can configure data collectors from HPE or third-party companies to forward events to OMi. Events are synchronized between servers. For example, OMi and OM synchronize the state of events and messages. If an OMi operator closes an event, a notification is automatically sent to OM. Similarly, OM notifies OMi about the acknowledgment of messages, and OMi automatically updates the lifecycle state of the corresponding events to “closed.”

Operators can enrich events with additional information, for example, by adding annotations to the event to either aid further problem resolution or to document what action has already been taken.

Closed events are automatically moved to the Closed Events browser. Operators can access this list of closed events, and can use these events as a reference for solving similar problems.

For those events that require the attention of specific subject matter experts, the operations bridge can forward those events to the appropriate operators. For example, the IT Operations System Administrator can configure the system to route notifications to operators and escalations to the appropriate help desk operators who can concentrate on managing escalated events and fixing underlying problems.

## Monitoring Dashboards

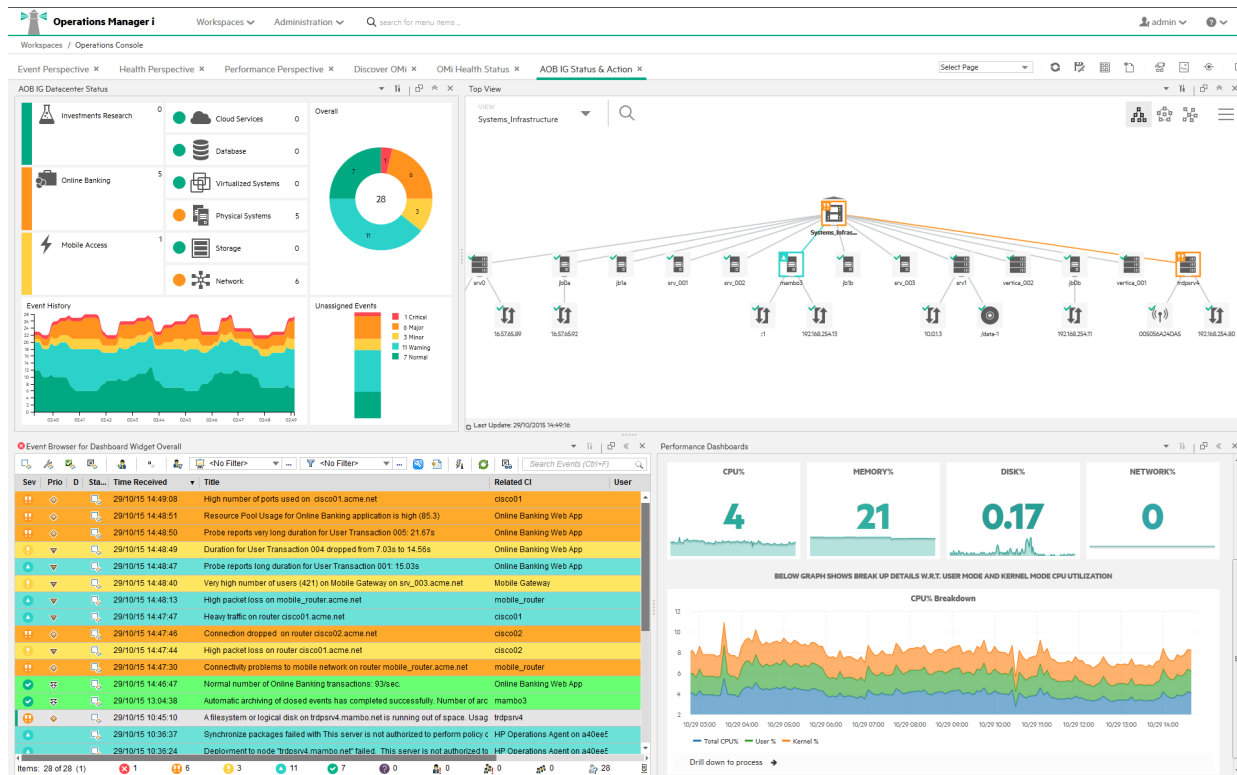
Monitoring Dashboards provide you an at-a-glance overview of the events from your monitored environment. They enable you to quickly assess the health of the environment and to identify areas that require your attention.

Monitoring Dashboards help you to:

- Get an overview of your monitored environment
- Visualize a starting point for daily management operations
- Quickly apply event filters to the event browser
- Keep an eye on the monitored environment while working on an event

Monitoring Dashboards display status information using widgets as building blocks (for example, stack and pie widgets). Each widget references an event filter, a view, or both, and only displays the status of those events that match the criteria of the filter and that are related to the configuration items included in the referenced view, making it easy to customize.

The following figure shows a Monitoring Dashboard screen:

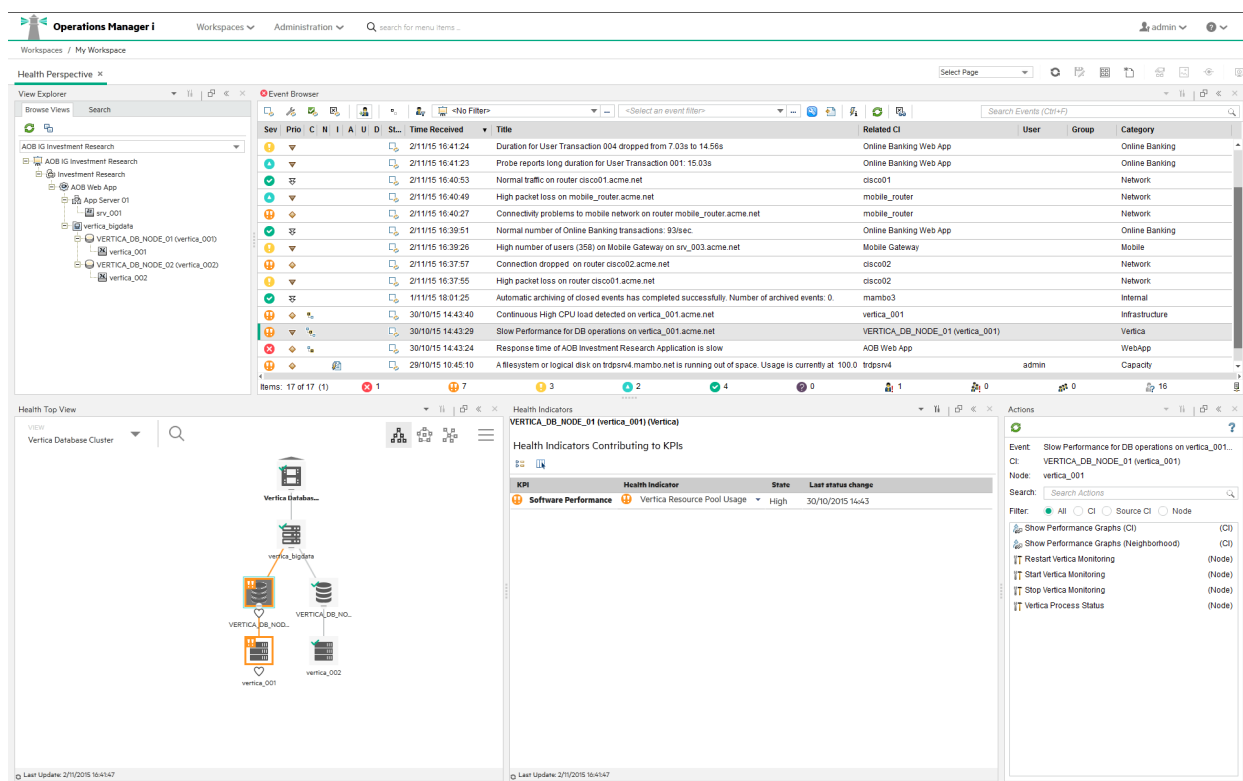


## Health Information

With event-based data, you can see in the Event Browser which related CI is affected by the event. Additionally, OMi health data, such as event type indicators (ETIs), health indicators (HIs), and key performance indicators (KPIs) are used to evaluate the health of related CIs in the context of the events.

For any device, such as a server, the severity of the problems directly associated with the server are collated and combined with information about devices associated with the server. The combined data is passed to calculation rules that evaluate and set the key performance indicators that indicate the overall health of the object.

The figure below shows a Health Perspective screen, with the Health Top View showing a hierarchical overview of the relationships among the objects associated with the event.



You can see the health status of an object, its KPIs, and Health Indicators, and see how the health status of the selected object affects the health of related objects. For example, you can navigate to check the health of neighboring CIs. This information helps you analyze which events to focus on, and prioritize event handling to maximize availability and minimize negative impact on business services. Users can also select views to show only the events and CIs they are responsible for.

## Correlating Events

In a large environment, one of the biggest challenges is how to manage the large number of events that originate from a variety of sources. Within this sea of data, the aim is to identify the events that have a significant impact on business services. So while it is essential to minimize the number of events that appear in the Event Browser, it is even more important to highlight the events that, if not managed properly, could cause a breach in service level agreements (SLAs) and generate incidents in your help desk system.

Event correlation plays a very important part in bringing together business service management and IT infrastructure management, where the disruption of a service can be traced to a specific failure in the IT infrastructure on which the service depends.

OMi correlates events automatically using the following forms of event correlation:

- Suppressing duplicate events
- Closing related events automatically
- Stream-based event correlation
- Topology-based event correlation

### **Suppressing Duplicate Events**

A new event may be a duplicate of an existing event. As a simple example, due to network stability problems, the same event is sent twice by the source domain manager because it did not receive an acknowledgment quickly enough for the first instance of the event. As new events are received, they are checked against existing events. If duplicates are found, new information, such as a change in severity, is used to update the existing event, and the new event is ignored. If duplicate event suppression is enabled, new events that are duplicates of an existing event are not retained and the original event is updated.

The advantage of correlating events using duplicate event suppression is that it reduces the number of events displayed in the console, but without losing any important information.

Suppressing duplicate events can result in additional correlations of the original event (both as cause or as symptom). When a duplicate is identified, the timestamp for the original event is updated to the time when the duplicate was received. The event is then correlated again and may now be related to other events which were not available for correlation when the original event was received.

### **Closing Related Events Automatically**

A new event can automatically close one or more existing events. When a new event arrives, a search is made for existing related events. Some specific information contained in the new event is used to match the new event to any existing events, and the new event closes the existing event. This type of event correlation is similar to the “good/bad message correlation” provided by HP Operations Manager.

For example, an existing event may be a notification of a problem or abnormal condition (a bad event) for a particular device. The bad event could be “SQL Query Performance SLOW”. Consider a new event matching this existing related event which notifies that the abnormal condition no longer exists (a good event). The good event could be “SQL Query Performance OK”. The new (good) event closes the existing (bad) related event.

You can track related events that were closed automatically in the event history.

### **Stream-Based Event Correlation**

Stream-based event correlation (SBEC) uses rules and filters to identify commonly occurring events or combinations of events and helps simplify the handling of such events by automatically identifying the events that can be withheld, removed, or need a new event to be generated and displayed to the operators.

The following types of SBEC rules can be configured:

- **Repetition Rules:** Frequent repetitions of the same event may indicate a problem that requires attention.
- **Combination Rules:** A combination of different events occurring together or in a particular order indicates an issue and requires special treatment.
- **Missing Recurrence Rules:** A regularly recurring event is missing, for example, a regular heartbeat event does not arrive when expected.

### Topology-Based Event Correlation

The event management process is simplified not only by consolidating events from all sources in a central console, but also by categorizing events using topology-based event correlation (TBEC). Dependencies between events are analyzed to determine whether some events can be explained by other events. For example, consider a database server (DB Server) running on a server (Server1). If the Server1 CPU usage becomes persistently overloaded, the resulting event “SLA for DB Server breached” can be explained by the causal event “Server1: CPU persistently overloaded (100% for more than 10 minutes)”.

The key is to pinpoint the underlying causal events that are responsible for other symptom events, so that you can prioritize the resolution of these causal events based on the impact to your business.

If two events occur concurrently (within a configurable time span), TBEC correlation rules identify one event as the cause and the other event to be the symptom. Rule-based event management enables you to manage large numbers of similar (related) symptom events in a large network.

When any combination of cause and symptom event occurs in the monitored environment, the correlated events are flagged in the Event Browser. You can configure the Event Browser to display the root-cause event and a separate overview of all the symptom events, thus enabling you to drill down into the correlation process and browse through the hierarchy of correlated events.

Events can also be correlated across technical domains, such as databases, hardware, networks, and web applications. This comprehensive scope enables you to correlate events that, at first sight, might not seem to have any connection. The cross-domain functionality also increases productivity by reducing the amount of overlap between operators responsible for monitoring different technical areas. For example, by correlating events relating to database problems, network problems, and storage problems, you can avoid the scenario of operators from the different technical areas all separately investigating different events that are the symptoms of one root cause event.

TBEC offers a number of benefits related to resolving complex events:

- Reduces the number of events displayed in the console, but without ignoring or losing important data that enables users to drill down through the hierarchy of related events.
- Supports event correlation across multiple domains to simplify root-cause analysis of events that generate symptom events.
- Changes to topological data do not require changes to correlation rules.

### Event Storm Suppression

If a problem is experienced on a managed system that results in the generation of an abnormally high number of events within a relatively short period of time, this phenomenon is known as an event storm. It is very probable that the root cause is already known and is being addressed. However, related events are also being generated. These events do not provide any useful information but may result in significantly increased loads

on the servers. To avoid this situation, OMi can be configured to look for event storms from managed systems and discard all subsequent events until the event storm condition for a particular system is over.

An event storm is detected when the number of events received within the detection time period, as a result of a problem on a system, exceeds the configured threshold required to enter an event storm condition.

When an event storm is detected on a system, events from this system are discarded until the rate of incoming events drops below the event storm end threshold. You can configure exception rules to select events from a system under event storm conditions that match a filter and either display these events in the Event Browser or close them (available in the Event Browser under Closed Event). The event storm end event automatically closes the associated event storm begin event.

## Structured Problem Solving

The centralized operations bridge streamlines the whole event management process. With centralized, consolidated information, you can create consistent, reusable, and optimized processes for event response.

You can deal with the majority of the events in your environment in a highly structured way. To help you manage events more efficiently and more effectively, you can use the following:

- **Tools**

You can create tools to help users perform common tasks on CIs. When you create a tool, it is associated with a CI type, and you can run the tool from the centralized console. For example, you can run a command tool to check the status of an Oracle Database instance. The tool is assigned to the configuration item type Oracle Database. If you are managing multiple versions of Oracle Databases, where the tool requires different parameters and options to check the status of the Oracle Database processes, you can create copies of the most appropriate tool and customize them for the various Oracle versions using the duplicate feature. Each tool is then dedicated to a specific version of Oracle.

- **Custom Actions**

You can automate your event management by creating actions to run on events to help solve problems and improve operator efficiency and productivity. Administrators can define a variety of custom actions for the operator to use when resolving certain types of events. Context-sensitive actions and context-specific tools can also be defined for specific circumstances. For example, you might create a set of database diagnostic tools that are designed to be used to help solve database problems.

For guidance about script definition and creation, including sample scripts provided with the product, see the *Operations Manager i Extensibility Guide*.

- **HPE Operations Agent Actions**

The events received in the Event Browser from HPE Operations Agent or OM may contain event-related actions configured corresponding policy templates in OMi or in OM policies. If event-related actions exist, you can run these actions from the OMi console. Those actions can be either operator-initiated or can run automatically when an event occurs.

- **HPE Operations Orchestration Run Books**

If you are already using HPE Operations Orchestration (OO) to automate operator tasks for analyzing or fixing problems, these OO Run Books can be mapped to CI types within OMi. You can launch OO Run Books in an event context from the OMi console.

In addition to manually launching Run Books, it is also possible to configure rules to automatically run a Run Book or a series of Run Books in the context of an event.

- **Graphs**



Graphs and charts provide additional data to help you visualize and analyze performance-related problems and trends affecting the CI impacted by an event, or any neighboring CIs. OMi graphs can display metrics from HPE Operations Agents, HPE SiteScope, HPvPV, BSM Connectors, and Application Management systems. Operators can even create their own personal graphs.

Structured event management processes are deployed to:

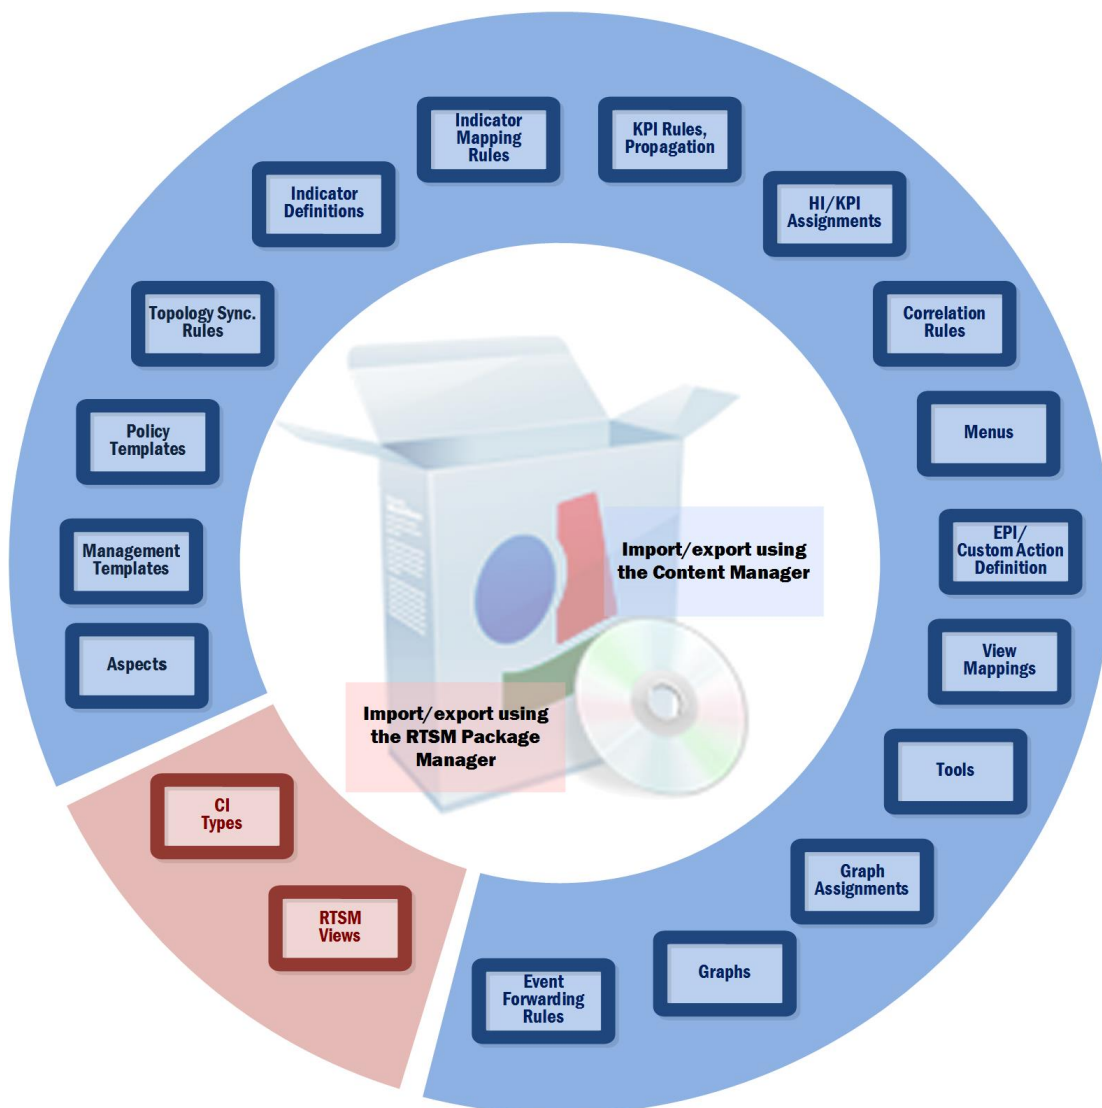
- Assign incoming events automatically to users in specific user groups. Automatic event assignment significantly increases the efficiency of event management and decreases the amount of time elapsed before a response to the event is possible. The IT Operations System Administrator can configure OMi to automatically assign incoming events immediately to available operator groups who are responsible for resolving those events.
- Start actions on events that match a specified set of criteria after a specified time. Time-based event automation rules consist of three main elements:
  - Filter defining the events to which time-based event automation rules are to be applied.
  - Time period defining the duration an event has to continuously match the rule filter to start the rule actions on that event.
  - List of actions to be started on matching events. Available actions are re-running automatic actions on events, modifying event attributes, forwarding events to external servers, assigning events to users and groups, running scripts, and running Run Books.
- Display and monitor the status of events using lifecycle management concepts. You can also see who is currently working on resolving the event, along with all other users who have already played a part in the solution.
- Document how an event is handled and solved. You can annotate the event to describe the problem resolution process, or capture domain expertise by tagging events with tips and hints that improve understanding and explain the event underlying problem.

## Management Packs

Management packs provide add-on content on top of OMi. They deliver automatic and end-to-end monitoring solutions of infrastructure and applications. Management packs enable users to monitor, detect, troubleshoot, and remediate issues in the IT domain. They increase the productivity of the user by optimizing and automating various tasks and reduce the mean time to resolve (MTTR) incidents.

Management packs discover application domains and proactively monitor the domains for availability and performance issues. They include, for example, management templates, aspects, policy templates, performances graphs, troubleshooting tools, auto remediation flows, Health Indicator and KPI definitions, as well as topology-based event correlation (TBEC) rules.

The following figure shows an overview of the content that can be included in a set of management packs:



### Out-of-the-Box Management Packs

A management pack provided by HPE typically consists of an RTSM package, a content pack, manuals, and the online help. All this content is automatically uploaded during the management pack installation. To use a management pack, a separate license may be required.

### Content Management Tools

OMi has a set of tools to help you manage your own content. These tools are the RTSM Package Manager and the Content Packs Manager. You can use them to package your own content and to exchange the content between systems. For example, you can prepare content in a test environment and then transfer the tested content to a production environment when the tests confirm that the content works as expected.

Export and import tools also enable you to exchange content between systems so that you can keep snapshots or backup images of the content you have developed. In addition, they make sure that different instances remain synchronized and up to date.

## Scalable Architecture with Multiple Servers

Operations Manager i enables you to manage widely distributed systems from a central location. In a distributed deployment, you can configure your environment hierarchically. You can then spread management responsibility across multiple management levels, according to criteria such as operator expertise, geographical location, and the time of day. This flexible management enables operators to focus on their specialized tasks, with the benefit of round-the-clock technical support available automatically and on demand.

The scalable architecture enables one or more OMi instances to be combined into a single, powerful management solution arranged to meet the requirements of your organizational structure. So you can configure servers to forward events to other servers in your environment.

In a distributed environment, servers hosting OMi can be configured to work not only with other like servers, but also with multiple OM for Windows and OM for UNIX management servers, other OMi servers, and third-party domain managers.

In such an hierarchical, distributed environment, you can configure OMi to:

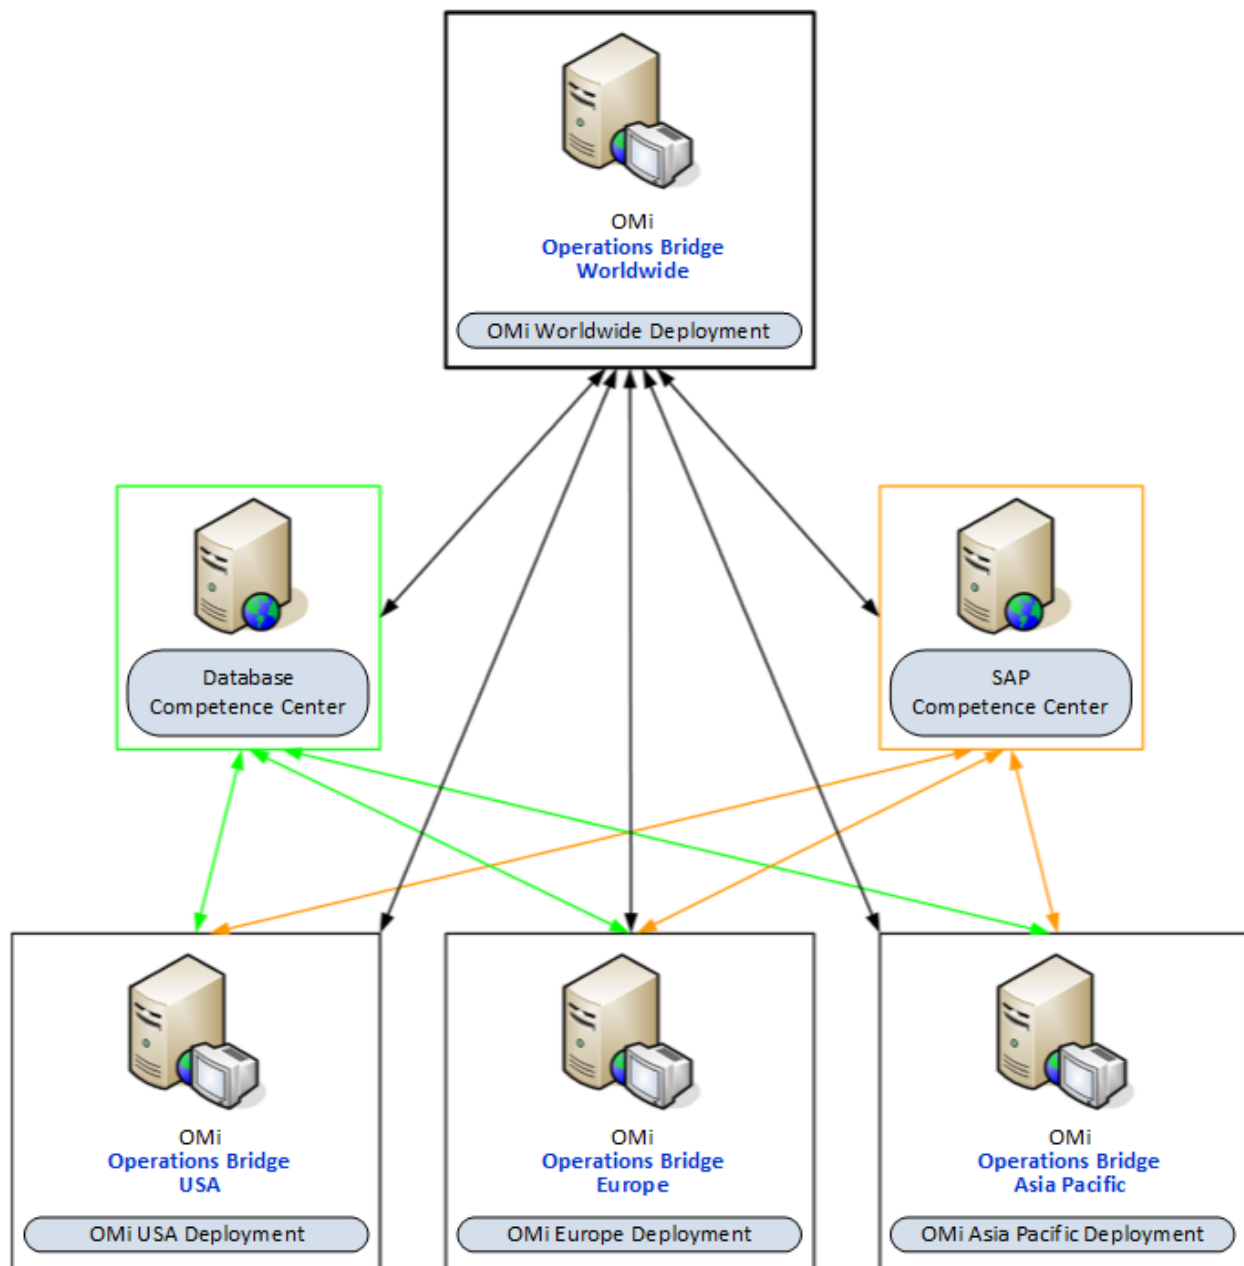
- Be the central event consolidator, or a manager-of-managers (MoM), for the whole environment at the top of the hierarchy.
- Work with other HPE products, such as NNMi and SiteScope.
- Work with third-party domain managers, such as Microsoft Systems Center Operations Manager.

You can configure servers hosting OMi to:

- Forward events to other servers hosting OMi and keep those events synchronized among the servers.
- Receive messages forwarded from multiple OM for Windows and OM for UNIX management servers and keep those messages synchronized between servers hosting OMi and OM management servers.
- Receive events forwarded from a BSM 9.x server receiving alerts from APM applications, such as HPE Business Process Monitor (BPM).

### Manager-of-Managers

The following figure shows an example of an hierarchical, distributed environment, with a central server hosting OMi managing other regional servers hosting OMi:



In this example, OMi Europe, OMi USA, and OMi Asia Pacific regional server deployments are managing different geographies. OMi hosted on the OMi Worldwide server deployment is at the top of the hierarchy, and is managing regional servers. It is acting as the central event consolidator, or MoM, for the complete environment. It is a worldwide operations bridge. Regional servers can also act as managers in their own geographies for subordinate systems to create a regionally monitored environment. It is possible to cascade the management of monitored environments in a hierarchical design.

If you operate in a large enterprise with multiple management servers distributed over a wide area, specialist knowledge relating to a specific subject is not always available locally. For example, your organization might have a competence center responsible for SAP. In addition, another center of expertise may be responsible for databases.

A competence center hierarchy distributes responsibility for configuration items in the monitored environment. Regional servers are not solely responsible for configuration items.

Instead, events about specific subjects go to a competence center server, where expertise exists to solve problems for similar types of configuration items in the monitored environment.

In a distributed environment, the IT Operations System Administrator can configure regional servers to forward certain messages to other servers in the network. The same System Administrator can configure regional servers to forward events to any server anywhere in the network, based on event attributes.

In the example scenario, all regional servers (OMi Europe, OMi USA, and OMi Asia Pacific) forward all database-related events to the database competence center server, and all SAP-related events to the SAP competence center server.

In this type of scenario, the operations bridge synchronizes event actions (for example resolve, assign, severity change) among the regional servers and the competence centers. This ensures the event states are always synchronized across the enterprise environment.

## Monitoring Automation


Monitoring Automation automates the configuration of monitoring for infrastructure and composite applications. Whether the monitoring tool used is agent-based or agentless, Monitoring Automation deploys the appropriate monitoring configuration to the target instances. It offers easy-to-tune monitoring and reporting, detecting new instances of components and mapping them to management templates and aspects that model the desired configuration and resource type.

Monitoring is the generation of events if a CI behaves in an unexpected manner. Typical events are:

- A monitored value exceeds a certain threshold. Example: Used disk space on a database exceeds a predefined limit of 90%.
- A node is removed from the network. Example: A power cut causes a server to shut down so it can no longer be reached.

Monitoring Automation provides a complete management solution for an application or service, enabling you to create a management solution for the entire set of configuration items (CIs) comprising the application. The solution can be made to respond dynamically to changes in the topology, making the monitoring solution independent of the hardware and platform running the application.

The key to understanding Monitoring Automation is to familiarize yourself with the underlying terminology and architecture. Consider the stack shown in the following figure:



The base of the stack represents the CIs to be monitored. CIs can be network elements such as computers, as well as applications or sets of applications providing a service. CIs are accessed in the following ways:

- Users interact with the CIs independent of any monitoring, as suggested in the central section of the figure.
- OMi monitors the CIs using the familiar monitoring structure shown in the left-hand section of the figure.
- A monitoring developer configures monitoring solutions as shown in the right-hand section of the figure.
- An application expert starts the monitoring process after tuning the configuration made by the monitoring developer, and acts on events passed by the operator by inspecting deployment jobs and using application-specific administration tools.

Monitoring Automation offers a number of features for creating flexible monitoring solutions. The following section explains each configuration element in turn. The explanation follows the order of the layers comprising the configuration stack from bottom to top.

### Node

A node is a physical element you can access on the network.

### CI

A CI is a node or an application or services running on a node. CIs are what is actually monitored by OMi. Events always relate back to CIs.

### Policy Templates

Policy templates define what is monitored and how the monitoring is done. Note that policy templates are platform-dependent.

Before Monitoring Automation, all configurations were done through Policies and Policy Templates, meaning that for each change in a CI with respect to the platform, the topology, or the monitoring policy, the values in the CI's policy templates against which the CI is monitored had to be modified.

### Parameters and Instrumentation

Monitoring Automation introduces parameters. Each parameter corresponds to a monitoring setting for a single CI attribute in the policy template. Changing the parameter value changes the monitoring behavior, removing the need to manually change hard-coded values in a policy template. The concept of cascading

default values is central to Monitoring Automation. The idea is that the monitoring developer or application expert uses as many default values as possible on a certain level, creating a baseline for monitoring. On the next level up, a subset of these values can and may need to be overridden for the specific monitoring task at hand, but every value already covered by the baseline setting can be taken over without having to redefine it.

The following features of parameters allow additional flexibility:

- Conditional parameter values enable using the same parameter with several policy templates, allowing hardware- and platform-independent monitoring solutions.
- Parameters with the same value can be combined into a single parameter. This removes the need to enter the same value multiple times.

Instrumentation includes scripts and programs executed by the HPE Operations Agent as defined in policies for managed nodes that have the agent installed on them.

### **Aspects**

Policy templates and instrumentation representing a certain expected behavior of the application or service to be monitored are grouped together in aspects. At aspect level, developers streamline the configuration as follows:

- They combine parameters with the same function into single parameters.
- They can nest aspects to combine aspects representing the same behavior, but defined in different policy templates, into a single aspect. Each nested aspect can be coupled with a deployment condition telling OMi which nested aspect is to be used in which environment. This allows any CI of the target CI type to use the same aspect, independent of the platform.
- They set default values at aspect level in line with the company monitoring policies.

### **Management Template**

A management template combines all aspects needed to monitor a composite application or service. The management template configuration includes the topology of the composite application and the aspects to be monitored. In addition, the developer overrides any company-wide default values at management template level if the application to be monitored requires this.

The developer hands the finished management template over to the application expert, who uses it to start monitoring the target application.

### **Tuning, Assignment, and Deployment**

Before starting the monitoring process, the application expert may want to override certain default values configured by the monitoring developer to take situation-specific monitoring requirements into account. This is called tuning.

The monitoring configuration represented by an aspect is defined in terms of a CI type. To start monitoring, this CI type has to be matched to an actual CI instance that has been discovered by the topology discovery process. This matching process is called assignment and can be done in the following ways:

- Manual assignment of a management template. The application expert links the management template to a CI instance of the management template's root CI.
- Manual assignment of an aspect. The application expert links the aspect to a CI instance of the aspect's target CI type.
- Auto-assignment. If the application expert defines auto-assignments for a management template or aspect, OMi dynamically assigns aspects to the relevant CI instances as and when they are discovered.

After assignment is completed, the monitoring solution is deployed in the same step. While the monitoring is running, the application expert can keep an eye on any deployment jobs to make sure the monitoring process proceeds as expected, or to acquire information related to events reported by an operator.

## User Engagement

The innovative User Engagement feature applies game dynamics to add extra stimulation to OMi users by providing business-enhancing challenges, accelerating operations bridge efficiency and user know-how. Successful progress through the various achievements is rewarded with Achievements and real-time notifications of great performance, helping to provide extra motivation to better engage with OMi, which improves users' performance in their daily work. Timelines are available to record each user's progress and collection of Achievements. Almost everyone is motivated by at least one of the types of challenges that game dynamics includes, for example, achievement, competition, status, and closure, and this makes User Engagement such a powerful feature.

By setting business-oriented achievements that OMi users work towards, and rewarding them for accomplishing the desired tasks, the most appropriate skills are being learned and the most important tasks are being completed while a level of engagement and excitement is being added to daily tasks. Users can watch as their efforts fill their achievement progress bars, and map their progress through their tasks and challenges in their dashboard. Completion of every new achievement can be accompanied with a popup notification providing immediate feedback of good performance.

User Engagement employs intrinsic motivations to help drive OMi users to achieve their set goals without the need to provide external benefits, which are generally accepted to only provide transient value. People naturally want to be successful and be seen to be successful. User Engagement provides the framework to help users learn how to use OMi and perform their daily tasks to a higher standard, being noticed for their achievements, so increasing the enjoyment and involvement in their work .

User Engagement administrators can select, configure, and enable built-in achievements tailored to the needs of their various OMi users. Users can work their way through their first-level achievements and once these have been successfully completed, they are invited to attempt the next level of achievements, increasing their perception of achievement and progress.

## Integration Interfaces

A number of interfaces are provided that enable integrations with other applications and allow modification and customization of the event management process. For example:

- To modify and enhance events during event processing, an event processing interface enables event processing scripts to be integrated into the event processing pipeline. This enables you to enrich events:
  - During event processing, for example, by adding information used in CI resolution and ETI resolution, or by influencing how duplicate events are handled.
  - To provide more information after event processing has taken place, for example, additional CI-related information from asset databases or information useful for troubleshooting purposes, such as a drill-down URL or links to external knowledge bases.
- To integrate events into other applications, an event web service interface enables developers and integrators to automate operator functions and event change detection. Most things that an operator can do



in the console while working on events can be done programmatically to improve efficiency. This interface also provides subscription support through Atom feed functionality.

- To synchronize events between OMi and an external event processing application, OMi provides an event synchronization web service interface. A typical use case is to synchronize events between OMi and an incident manager, such as Service Manager.
- To integrate directly with other domain managers, such as Microsoft Systems Center Operations Manager, OMi provides the HP BSM Connector.

The OMi Extensibility Guide in the OMi documentation library describes these interfaces and provides information for content developers and integrators to customize and extend the functionality of OMi.

## Business Value Dashboard

OMiBusiness Value Dashboard brings your OMi data to life. Use BVD to create custom, flexible dashboards that visualize information from OMi and other sources in an informative and appealing way. Your BVD dashboards can be accessed anywhere, anytime, from any device. Incorporate your own graphics, add color to identify status, and receive real-time updates—so you always understand the value driven by your IT environment.



**Anytime, Anywhere.** BVD dashboards are real-time dashboards. You choose how often you send data to BVD; BVD displays the data with no delay. You decide where you want to view your dashboards: PC, tablet, or phone. BVD supports the major browsers. Choose your favorite!



started.

**Simple, Colorful, Flexible.** Design your dashboards using Microsoft Visio. BVD provides a Visio stencil with shapes that then later become the widgets in your dashboards. The shapes include widgets for drawing charts, coloring text or values, displaying information feeds, web pages (for example, video streams), and many more. BVD provides sample dashboards to help you get



**Connect.** Once uploaded to BVD, you connect your widgets to the data. The BVD Manage Dashboards page makes this task simple and efficient. You can set additional widget options such as rules that determine the visibility and status colors of the widgets; or you can link widgets to other dashboards to enable drill down.



**Integrate.** BVD can process any kind of data as long as it is sent in JavaScript Object Notation (JSON), a language-independent, open data format. The out-of-the-box integration with OMi facilitates the integration of event and KPI status as well as metrics data. BSM Connector provides policies that automatically forward data collected from various sources to BVD.

Alternatively, create your own integrations for any data source by writing an adapter for BVD. The adapter must convert the source data to JSON and send the JSON-enabled data to the BVD data receiver.

Out-of-the-box, you can easily configure OMi to send the following data to BVD:

- **Event status data:** The event status to be forwarded is collected from an OMi monitoring dashboard that you specify. Use the `bvd-event-status` command-line interface on the OMi server to forward event status.
- **KPI status data:** The KPI status is collected from all CIs that are associated with a view that you specify and that have the KPI set that you specify. Use the `bvd-kpi-status` command-line interface on the OMi server to forward KPI status.
- **Metrics data:** The metrics data is collected from your graph favorites in OMi. To forward metrics data, enable data forwarding in Performance Graphing (known as Performance Dashboard in OMi 10.10), then save your graphs as favorites with the export (or forward) data option selected.

BVD is part of the OMi package, but comes with its own installer. You can install BVD on a gateway server, or on a separate server. For more detailed information on BVD and specific instructions on the integration with OMi, see the BVD Help.

## User Roles and Responsibilities

Installing, configuring, and running the operations bridge requires a team of people who have special skills and domain expertise. Each role has a different set of responsibilities and tasks.

- The Operator is the hands-on event manager and troubleshooter.
- The Monitoring Developer knows both the monitoring product and the application well enough to be able to develop the monitoring solution. He decides what is to be monitored and what the appropriate performance levels should be.
- The IT Operations System Administrator installs and configures the monitoring and event management processes. What he can configure is very flexible. He adds new users in the OMi area according to local requirements. He can grant permissions and restrict access to Administrative UIs, Tool Categories, and Custom Actions. He can specify rights and permissions for individual users or user types. He can also enable or disable access to events assigned to other users. For example, he can enable users to view events that are not assigned to them, but deny them the right to make any changes.
- The Application Expert knows everything about a specific application or service. She administers the equipment involved in running the application and troubleshoots it if monitoring events indicate there is a problem.

Frequently encountered titles for these user roles, together with a summary of their responsibilities, are presented in the table below. Now that we know more about OMi, we will follow some typical users in subsequent chapters to see how they manage their workday and complete their tasks. In the next chapter, we learn more about the daily responsibilities of Dave the operator in an enterprise environment with OMi as the operations bridge.

Job Title	Other Titles	Responsibilities
<p>Operator</p>  <p>“Dave”</p>	<ul style="list-style-type: none"> <li>• Domain Operator</li> <li>• IT Operations Operator</li> </ul>	<p>Monitors daily events assigned to him or his workgroup.</p> <p>Performs routine non-OMi operations on the applications, systems, networks he is responsible for.</p> <p>Troubleshoots and resolves events that might escalate into an incident.</p>
<p>Monitoring Developer</p>  <p>“Mike”</p>	<ul style="list-style-type: none"> <li>• Domain Expert</li> <li>• IT Operations Monitoring Developer</li> <li>• Subject-Matter Expert for applications, networks, or other specialized areas</li> </ul>	<ul style="list-style-type: none"> <li>• Customizes the way OMi monitors a domain.</li> <li>• Configures management templates, aspects, and policy templates for Monitoring Automation.</li> </ul>
<p>Administrator</p>  <p>“Matthew”</p>	<ul style="list-style-type: none"> <li>• System Administrator</li> <li>• IT Operations System Administrator</li> <li>• OMi Administrator</li> <li>• System Architect</li> </ul>	<p>Oversees the OMi environment and task assignments.</p> <p>Integrates OMi with other tools and processes.</p>
<p>Application Expert</p>  <p>“Alice”</p>	<ul style="list-style-type: none"> <li>• Subject Matter Expert for a certain application or service</li> <li>• Application Administrator</li> </ul>	<p>Tunes a monitoring solution to the specific environment of her application or service and assigns management templates or aspects to system nodes.</p> <p>Deploys the monitoring solution and assures monitoring is running correctly.</p>

# Chapter 3: Operator Workflow



We met Dave in chapter "Introduction to Operations Manager i". Dave is the operator responsible for daily event management in an OMi deployment. An operator is usually an entry-level position in the corporate IT environment, but Dave has diverse skills that he brings to the position because he has experience with many of the technologies in the Operations Manager i environment.

Dave works a varied schedule because he is often called when problems occur. He may solve them in person, or log in remotely to ensure that his user community can work without interruption. The Operations Manager i user interface enables him to monitor the events in his domain from any location as long as he has network access.

Dave needs to understand event management and how to use all the health-related tools at his disposal. There are tools, self-configured commands, scripts, and links to other information that help operators like Dave resolve and close different types of events that occur in the operational environment.

The operations bridge enables Dave to see alerts and events in his domain immediately. He can concentrate on managing his events and fixing the underlying problems automatically with the

appropriate tools.

Dave adds value to the enterprise by prioritizing the events in his domain according to their impact on business services and continuity. Dave must resolve small problems before they become major problems that lead to degradation in the quality of supported business services.

Experience with underlying technologies can help Dave correlate events that occur in different technical domains, such as: databases, hardware, network, web applications, and so on. He monitors these disparate technologies to minimize the impact of a failure in one area that might reduce system responsiveness in another area. Minimizing problems before they escalate improves enterprise productivity by minimizing the cascading effect of an unidentified critical event.

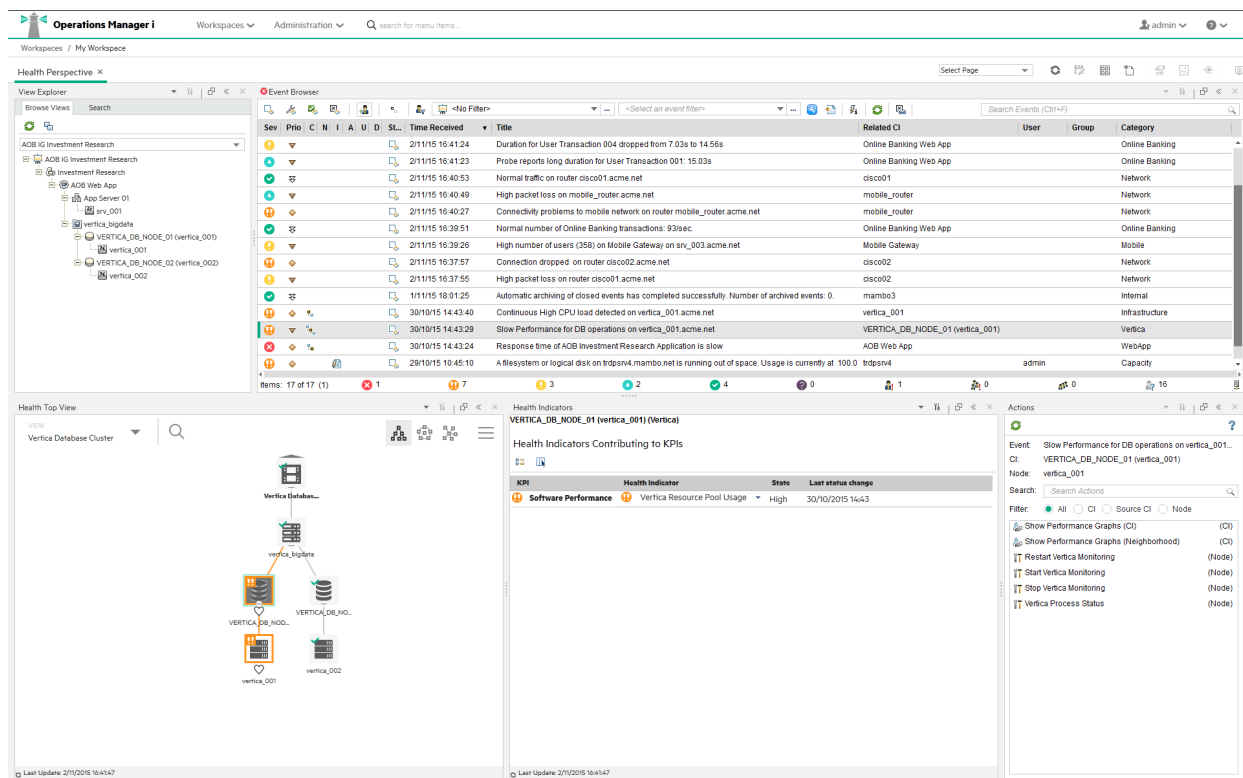
If Dave cannot fix a problem, he can escalate it by forwarding the event to an external event processing application. This usually including transferring ownership of the event, for example to a help desk operator or an application expert.

## The Operator Environment

The system administrator determines the events that each operator can view or modify by defining user roles and assigning user rights. Dave can see his assigned events, plus other events that he is allowed to see, in a cross-domain view. For example, he is responsible for maintaining the enterprise e-mail server, but he might be able to see events that are assigned to another operator.

## Health Perspective

The following figure shows the Health Perspective tab with five panes that show different views of the system. Dave begins every day by opening the Health Perspective:



The five panes provide a global view of the events in Dave's domain:

- The Model Explorer enables Dave to select a view and an area that he is responsible for. The view shows the parent child relationships among the CIs.
- The Event Browser lists all related events and related information in a table view.
- The Health Top View of a selected event shows the key performance indicators (Kips) of the CI related to the event, and the CIs in its neighborhood.
- The Health Indicators pane provides detailed information about the status of any CI selected in the Health Top View pane. This view shows information about the performance, availability Kips, and any health indicators that are relevant to the selected CI.
- The Actions pane is used to display the actions that are available for the selected event, its related CI, or the node that hosts the CI. Actions include tools, run books, custom actions, and performance graphs.

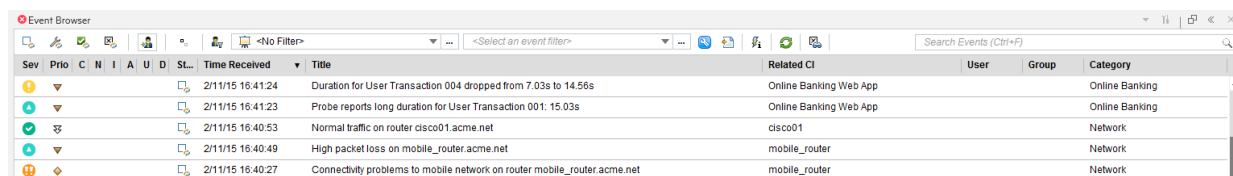
## Event Browser

The Event Browser is the first area Dave looks at. He can see:

- A list of prioritized active events.
- Events assigned to him.
- Information about unresolved and unassigned events.

- Tab details that show how many events are critical, major, minor, warnings, normal, or the status is unknown.

The following figure shows a typical global view of event information arranged in the Event Browser pane:




Sev	Prio	C	N	I	A	U	D	SL	Time Received	Title	Related CI	User	Group	Category
!									2/11/15 16:41:24	Duration for User Transaction 004 dropped from 7.03s to 14.56s	Online Banking Web App			Online Banking
!									2/11/15 16:41:23	Probe reports long duration for User Transaction 001: 15.03s	Online Banking Web App			Online Banking
!									2/11/15 16:40:53	Normal traffic on router cisco01.acme.net	cisco01			Network
!									2/11/15 16:40:49	High packet loss on mobile_router.acme.net	mobile_router			Network
!									2/11/15 16:40:27	Connectivity problems to mobile network on router mobile_router.acme.net	mobile_router			Network

Dave uses filters to see events from out-of-box views, or he can personalize his workspace by customizing filters and tabs. For example, he can use a combination of severity and priority to identify the events that need immediate attention. The first task is to determine which of the highest priority events should be examined first.

### Health Top View

When Dave selects an event to investigate, the Health Top View is updated to show more information about the related CI. For example, assume that the event is caused by an exceeded storage quota on a related server. The Health Top View shows the topological view of the affected server. Dave can select it in this view to obtain more information. The following figure shows a typical Health Top View of business services and CIs:

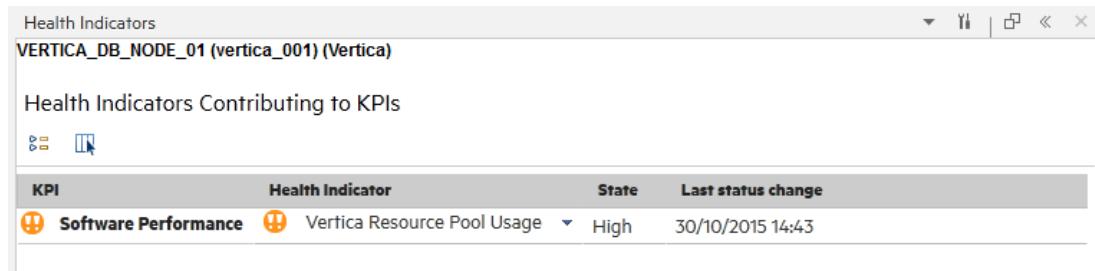


The Health Top View enables an operator to see the health of any CI in the topological diagram. Examining upstream and downstream CIs may provide other clues that help to isolate the problem.

The next step is root cause analysis using the Health Indicators pane.

### Health Indicators Pane

When Dave selects the affected CI in the Health Top View, the Health Indicators pane is updated to display more information about the CI.



This detailed view can show whether there is one underlying problem or a variety of contributing factors to the critical event. This information enables Dave to make quicker decisions about what he needs to do next. As he takes action, other operators will see that Dave is working on this problem so that they can concentrate on other critical events.

Dave may also use performance graphs and other tools to troubleshoot the problem.

### Other Tools

The details of an event can contain instructions. Dave can select the **Additional Information** tab, which might contain notes or other tips to solve the problem. There may be a diagnostic tool or script that he can run to analyze CI performance in great detail, or related logs with informative error messages.

Dave has performance graphs at his disposal that are useful analysis tools. For example, if a database performance event occurs, Dave can right-click the event and select **Show > Performance Graphs (Neighborhood)**. Performance graphs are displayed for the CI affected by the event and for its neighbor CIs, such as the affected application server. These graphs show not only the performance information at the time of the event, but can also show performance at an earlier point in time.

**Note:** Operations Manager i tools are not limited to troubleshooting events. Dave can also launch tools just to perform routine daily tasks.

### Resolution

There are many ways to solve a problem. For this example, Dave sees a suggestion to run a tool from the **Launch** menu. From the Event Browser, Dave right-clicks the event and selects **Launch > Tools > Repair File System (CI)**. When the tool finishes, the problem is resolved and the event disappears from the list. If this did not work, Dave can access related run books from the Actions pane. Run books are scripts that execute a multi-step process to solve the problem.


## Other Roles

The operator relies on the expertise of two other key roles:

- The monitoring developer.  
See ["Monitoring Developer Workflow" on page 33](#)
- The system administrator.  
See ["IT Operations System Administrator Workflow" on page 36](#)



# Chapter 4: Monitoring Developer Workflow



Mike is an OMI monitoring developer. His primary focus is to customize OMI to meet specific business requirements.

Typically, Mike integrates new applications and CIs into the monitoring process. To create a monitoring environment for composite applications and services in line with company standards, Mike configures Monitoring Automation elements such as management templates and aspects, and implements the technical monitoring details by customizing policy templates and linking them to management templates and aspects. He also makes monitoring solutions independent of instances and platforms by using auto-assignments and conditional deployment of aspects.

For example, he must define how to monitor a new group of servers that support FTP. These servers support critical business services by enabling internal departments and services to send and receive large data blocks.

Initially, the payroll department will transmit employee payroll information to the corporate payroll service from this server. The payroll service will send back internal summary reports, required governmental reports that must be archived and re-distributed to

the appropriate government agencies, and deliver other payroll related outputs.

Mike must define how to monitor FTP servers to ensure their ongoing health, ensure they can process requests, and permit FTP downloads in a secure environment. If the servers are unavailable, deadlines will be missed, and in extreme cases an outage could generate fines for failing to meet governmental requirements.

## Initial Analysis

The first things Mike needs to think about are the KPIs and health indicators for FTP servers. Some key questions to answer:

- If application availability and performance is important, how should he measure those KPIs?
- What are the service level agreements (SLAs) that might be breached if FTP server availability and performance thresholds are not met?

The IT organization is not only responsible for ensuring this server and its resident applications are available, but also that all associated resources perform according to user expectations. Mike needs to factor all of this information into his selection of KPIs and definition of key health indicators.

## Define Health Indicators

What are the KPIs that should be measured by the monitoring process?

How should they be reported back to business service owners who use the FTP server?

Mike turns his attention to defining the health indicators that support the selected KPIs. For example, health indicators for the application availability of a Windows FTP server could include Windows service metrics that report:

- Number of outbound connections of all types from the service
- Number of transferred bytes per second
- Server response time

Mike must configure the monitoring process, create the monitor policy, and determine how to report its status.

Fortunately, he has several HPE applications that support these tasks. His experience and overall knowledge of these applications help him select the best fit for the task. For example, he might choose an Operations Manager agent policy, SiteScope, or another HPE monitoring tool. Whatever he selects as a health indicator must have a supporting tool that can report the status of the health indicator.

## Configure Monitoring Automation

Finally, Mike needs to think about how he wants to implement the monitoring process and automation in OMi:

- What does he need to monitor and how should the monitoring process be configured?
- How can he configure policy templates, parameters, instrumentation, aspects and management templates to monitor the Health Indicators defined previously?

Mike decides to create a management solution for the entire set of CIs associated with the FTP server. For this, he considers several features of monitoring automation:

He uses policy templates to define the details of the monitoring tasks, and parameters and instrumentation to increase flexibility. Mike then creates aspects to streamline the configuration. He also uses a management template to group all aspects used in the FTP server monitoring process and assigns the monitoring configuration to the specific CIs.

Mike considers making his monitoring solution independent by using auto-assignments and conditional deployment of aspects.

## Other Tasks

There are a variety of tasks that Mike completes to enrich the monitoring and health maintenance process for the FTP server. He might do one or more of the following:

- Create graphs that summarize the metrics collected for the FTP server, and assign them to the FTP server CI type to make them appear automatically.
- Create OMi tools to restart the FTP server.
- Create multiple operational run books. For example, Mike can create a run book to delete obsolete files from the FTP server.
- Create content packs that contain the monitoring artifacts.
- Create correlation rules to map certain identified disk problems to certain FTP server problems.

Mike has an important role. He envisions what metrics are necessary, how they will be captured, and defines the related processes to gather data and solve problems.

## Other Roles

Mike, the monitoring developer, integrates new applications and CIs into the monitoring process. These are configured by Matthew, the IT Operations system administrator, for use by the operators, Dave, and his colleagues. He also develops management solutions for use by Alice, the application expert, and her colleagues.

For an insight into these other personas, see:

- The system administrator.  
See ["IT Operations System Administrator Workflow" on page 36](#)
- The operator.  
See ["Operator Workflow" on page 28](#)
- The application expert.  
See ["Application Expert Workflow" on page 40](#)

# Chapter 5: IT Operations System Administrator Workflow



In chapter "Introduction to Operations Manager i", we learned about the concept of an operations bridge. OMi is the operations bridge for a complete Business Service Management solution, providing a centralized location for event and performance management.

In chapter "Operator Workflow", we learned that the operations bridge provides a complete view of all operational events to enable an immediate response whenever necessary. To run efficiently, someone must configure and optimize the operations bridge. That is Matthew's task as the IT Operations system administrator.

Matthew is behind the scenes, designing an efficient monitoring environment for the operations staff. In his role, he ensures ongoing maintenance, manages users and user roles, and looks for opportunities to fine-tune the monitoring process. He designs the operational system and puts the processes in place for others to use on a daily basis. Creating new scripts and automating as many processes as possible is his specialty.

Matthew must have in-depth knowledge of the operational environment, understand the dependencies among applications, and configure an environment that is as efficient as possible.

## Installation and Configuration Tasks

Matthew has the global expertise to install, configure, and integrate OMi with other applications, such as HPE Operations Orchestration or HPE Service Manager, and configures event forwarding from various sources, such as HPE Network Node Manager i (NNMi) or OM systems.

Matthew also enables the monitoring process by installing the required monitoring tools, such as Operations Agents and SiteScope.

Matthew also installs and maintains management packs or custom content packs.

If required, Matthew installs the HP BSM Connector to assist with the integration of third-party domain managers, such as Microsoft System Center Operations Manager.

Matthew has these responsibilities:

- [Oversee the OMi Installation](#)
- [Tune the Environment](#)
- [Tune Infrastructure Settings](#)
- [Configure Users and User Roles](#)

### **Oversee the OMi Installation**

Matthew has domain expertise and experience with OMi. He understands how to install OMi and how to configure it correctly. He designs and supervises the end-to-end installation process of required OMi components and decides which applications should integrate with OMi. These applications include other HPE solutions and third party applications, such as Microsoft SCOM.

The complexity comes from integrating multiple infrastructure and enterprise business applications according to Information Technology Infrastructure Library (ITIL®) principles. The goal is to set up and configure autonomous applications that work seamlessly with one another. Each operates independently but communicates effectively with other applications.

### **Tune the Environment**

Matthew configures all the connected servers. Then he sets up rules for forwarding events and notifications and decides who should receive the event notification. In some cases, the event response is to use the custom scripts that Matthew identifies, or even produces himself. Finally, he designs the process that assigns new events to a specific user group. These are rule-based filters to ensure that OMi automatically assigns each event to the right group or individual.

### **Tune Infrastructure Settings**

These settings represent a large area of required expertise. If Matthew changes a setting, he has to understand the resulting impact on the operational environment. For example, if he limits what is written to the audit log, details of certain events will be omitted. Other settings describe different aspects of the environment (such as the SSL certificate server), how related events are managed, and duplicate event management.

### **Configure Users and User Roles**

Matthew is responsible for defining user roles and the rights and limitations that accompany these roles. The user role is a generic way to assign the same rights to users, instead of configuring each permission separately. If a new operator or monitoring developer joins the staff, Matthew adds them to the system and assigns one of his pre-defined user roles to automatically grant the same rights and limitations that everyone else with that user role has.

## Other Responsibilities

Other responsibilities include:

- Deciding which Event Processing Interface (EPI) scripts to run at pre-defined times
- Defining custom actions
- Defining Workspace pages and Monitoring Dashboards for different users

## Ongoing Tasks

After initial installation and configuration, the beneficiaries are the operators whose task is to manage the events they monitor. Mike delivers an environment to Dave the operator that simplifies his daily tasks and ensures that he can respond to critical events as quickly and efficiently as possible.

After initial configuration, maintenance is automatic until a user requires a change. Most environments must also change over time to meet new demands. Mike the monitoring developer might send new or updated content packs for Mike to install. As the enterprise grows, Mike must add new users and assign each one the appropriate user role and permissions. Mike may also deploy patches for Operations Agents when needed.

Mike also can see from daily operations that he needs to revise some of his original models for event forwarding and notifications. As new situations present themselves, Mike decides whether to use existing scripts or create new response models. Tuning the environment makes the operation more efficient and monitoring more effective.

## Operations Bridge

By gathering all infrastructure operations, including applications, dedicated servers, and related software and hardware under a single IT umbrella, it is possible to meet enterprise service level objectives. Mike's role is to configure this high-performance environment and use OMi as the operations bridge. All components work in concert to deliver necessary internal business services to employees, and provide portal services or other application availability to external customers. Imagine an international banking environment with arrays of servers, applications, CIs, and more to ensure a 99.999% response. This type of commitment requires the type of well-designed operational environment that Mike provides.

## Other Roles

Matthew, the IT Operations system administrator configures and optimizes the operations bridge, including content developed by Mike, the monitoring developer, for use by the operators, Dave and his colleagues.

For an insight into these other personas, see:

- The monitoring developer.  
See ["Monitoring Developer Workflow" on page 33](#)
- The operator.  
See ["Operator Workflow" on page 28](#)

# Chapter 6: Application Expert Workflow



In chapter "Introduction to Operations Manager i", we learned about the concept of an operations bridge. OMi is the operations bridge for a complete Business Service Management solution, providing a centralized location for event and performance management.

We also saw how Monitoring Automation can help to create a flexible monitoring solutions for applications and services.

In chapter "Monitoring Developer Workflow" we met Mike, who designs monitoring solutions in line with the company's policies on what should be monitored and how it should be monitored.

Alice is the application expert for a particular application or service, and is the person who knows most about the systems the applications runs on, and how the application is used. Alice is in charge of deploying the management template developed by Mike to monitor the actual application instance for which she is responsible.

## Installation and Configuration Tasks

Before starting to monitor her system, Alice tunes the values against which the application is to be monitored. The values configured into the management template by Mike, the monitoring developer, reflect the company-wide standards for monitoring applications of the type of Alice's application. Alice may need to change some of those values to suit the particular application instance she is responsible for. She decides whether values need to be changed in the management template or the automatic assignment rule or whether she will overwrite values for specific CIs representing the particular application instance manually.

- OMi discovers instances of the configuration item types in the topology view configured into the management template. All Alice needs to do is to define the auto-assignments that can be done for the management template. After finishing the configuration of the auto-assignments for the management template, OMi matches the configuration item types in the management template to discovered configuration item instances, and deploys the aspects required to monitor them automatically.
- If more control is required, Alice can manually assign the management template or aspect to discovered configuration items, after which OMi deploys the aspects in the management template.

## Ongoing Tasks

Alice may be contacted by operators, for example Dave, if multiple events are generated even though no real problem occurred. This may be caused by incorrect thresholds that are too low. In this case, Alice will fine-



tune the monitoring configuration and change threshold parameters for certain CIs, in automatic assignment rules or inside management templates.

## Other Roles

Alice, the application expert, tunes and initiates the monitoring process for the application instance she is responsible for using a management template developed by Mike, the monitoring developer. The monitoring process generates events that are handled by Dave, the operator.

For an insight into these other personas, see:

- The monitoring developer.  
See ["Monitoring Developer Workflow" on page 33](#)
- The operator.  
See ["Operator Workflow" on page 28](#)

# Summary

After reading about the different users who install, configure, and manage the day-to-day operations of Operations Manager i, you can see that it takes multiple skill sets to make everything run at optimum level. You may fill one of the roles described in this guide. Regardless of which role you assume, you can make a difference in how well your work group delivers value to your internal customers.

# Send Documentation Feedback

If you have comments about this document, you can [contact the documentation team](#) by email. If an email client is configured on this system, click the link above and an email window opens with the following information in the subject line:

**Feedback on OMi Concepts Guide (Operations Manager i 10.10)**

Just add your feedback to the email and click send.

If no email client is available, copy the information above to a new message in a web mail client, and send your feedback to [ovdoc-asm@hpe.com](mailto:ovdoc-asm@hpe.com).

We appreciate your feedback!



Go OMi!