IA L06 – No-fail Failover with Disaster Recovery Advisor

“Become proactive in managing HA / DR”

Hands-on Lab

Description	Get hands-on experience with the latest version of Disaster Recovery Advisor (DRA) as we explore the various functionalities within the product. Visualize how DRA can play a critical role in drastically reducing the time and effort associated in managing your disaster recovery preparedness. Utilize the analytics and reporting tools to assess High-Availability (HA)/Disaster Recovery (DR) risks. Walk away with the know-how to generate actionable remediation reports on those HA/DR risks within your environment. This lab will provide you with a detailed look into some of the functionality behind DRA. You get a better understanding of the products potential by navigating and controlling the user interface as we walk through this hands-on exercise.

At the end of this lab, you should be able to	▪ Use DRA to help minimize service availability risks in the datacenter. ▪ Learn how to identify the root cause of production, cloud and DR service availability issues. ▪ Asses your cloud infrastructure reliability. ▪ Improve cross-team coordination in achieving reliable and redundant IT services. ▪ Maintain compliance with service availability goals. ▪ Identify storage optimization opportunities.

Notes	
	▪ A brief presentation will introduce this lab session and discuss key concepts. ▪ The lab will be directed and provide you with step-by-step walkthroughs of key features. ▪ Feel free to follow the lab using the instructions on the following pages. You can optionally perform this lab at your own pace. ▪ Be sure to ask your instructor any questions you may have. ▪ Thank you for coming to our lab session.

Topics and estimated times:

Primary Exercises:

1) **Proactive Risk Detection**

   ≈ 10 minutes

2) **Configuring Comparison Analysis**

   ≈ 10 minutes

3) **Proactive Risk Detection in the Private Cloud**

   ≈ 10 minutes

4) **Detect Recovery Point Objective (RPO) Service Level Agreement (SLA) violations**

   ≈ 10 minutes

5) **Cross-domain Collaboration**

   ≈ 10 minutes

Extended Exercises:

1) **Create customized, automated email notifications**

   ≈ 5 minutes

2) **Schedule high-level customized reporting**

   ≈ 8 minutes

Login to the Management Server Console

Action: To begin the exercises, open the **Disaster Recovery Advisor** link on the desktop and login:

1) Enter Username: **admin**
2) Enter Password: **symantec**
3) Click **Login.**

(*) Note: password is case sensitive (use all lowercase)

Lab Exercise 1

Topic – Proactive Risk Detection	Duration ≈ 10 minutes

The scenario for the first exercise is the following:

- Our Enterprise resource planning (ERP) system has recently undergone an upgrade.
- Prior to owning DRA, there was no simple way to assess readiness and measure risk following IT configuration changes.
- DRA has identified issues in configuration following the upgrade and provided us with open tickets to remediate the problems.

Get immediate reliability status

- DRA has identified open risks in the areas that are red and yellow.
- These risks are grouped by business entity and category.
- For example, there are a total of 9 data protection risks in the ERP business entity.
- Let's investigate those 9 risks by clicking on the red indicator under **Data Risk** and next to **ERP**.
- This will take you to the internal ticketing interface.

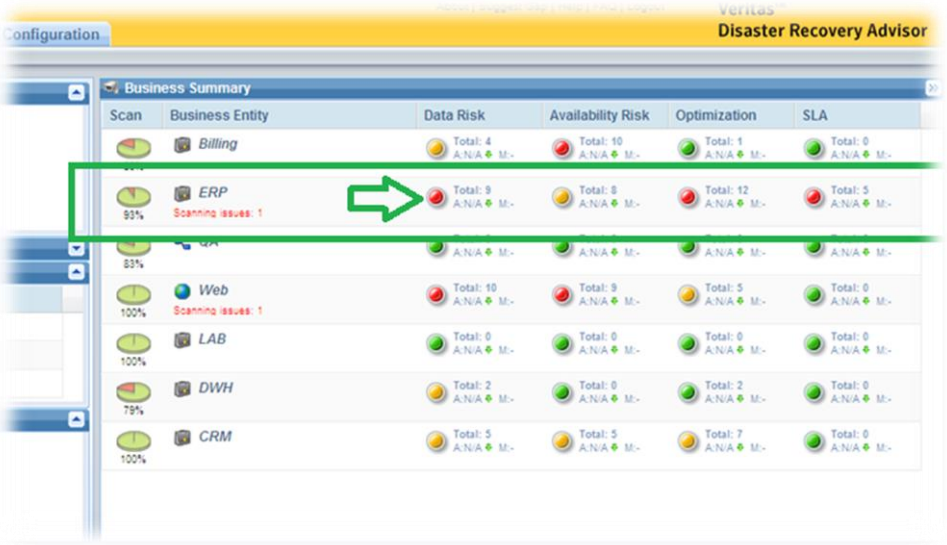

Detailed information about faults and risks is now presented which is filtered by business entity (**ERP**) and category (**Data Risk**). Let's refine the filter to get a better view of the open tickets in this area.

Change the default grouping from "**Gap Type**" to "**(no grouping)**."

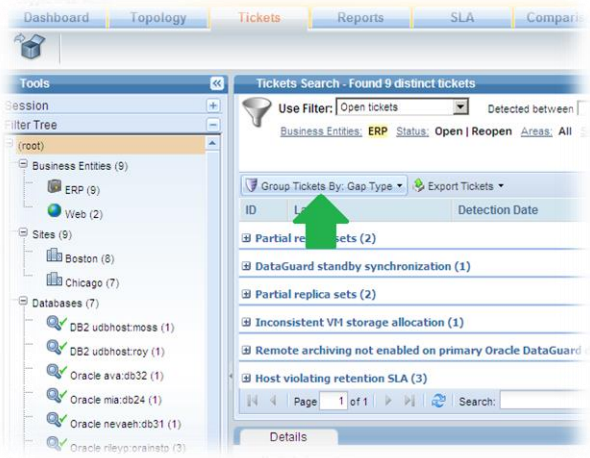

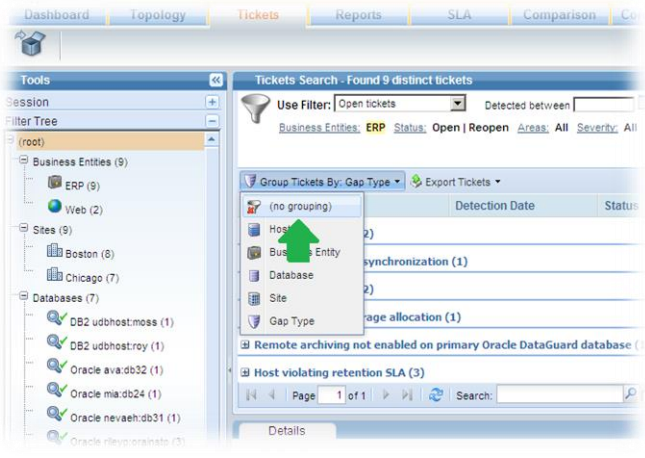

Now sort the list of identified risks by priority. Click on the "**Rating**" column header until a blue triangle pointing down appears. The list is now sorted with highest priority items first.

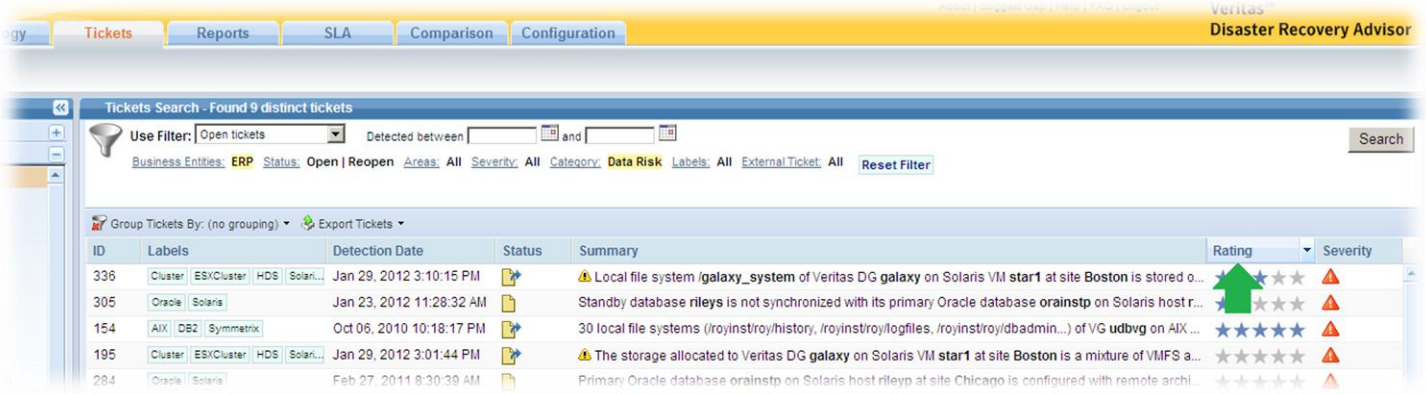

Each identified risk is presented in a DRA Ticket. As we'll see later, you can easily integrate the fully-functional, built-in ticketing engine with any existing ticketing or incident management tool.

The ticket summary table captures the high-level details of each risk, including the detection date, a brief description, the technical severity, the priority to fix ("rating"), labels (useful for compound searches), and more.

To learn more about a ticket, click on it. Let's click on the ticket with the highest priority (**ID 154**):

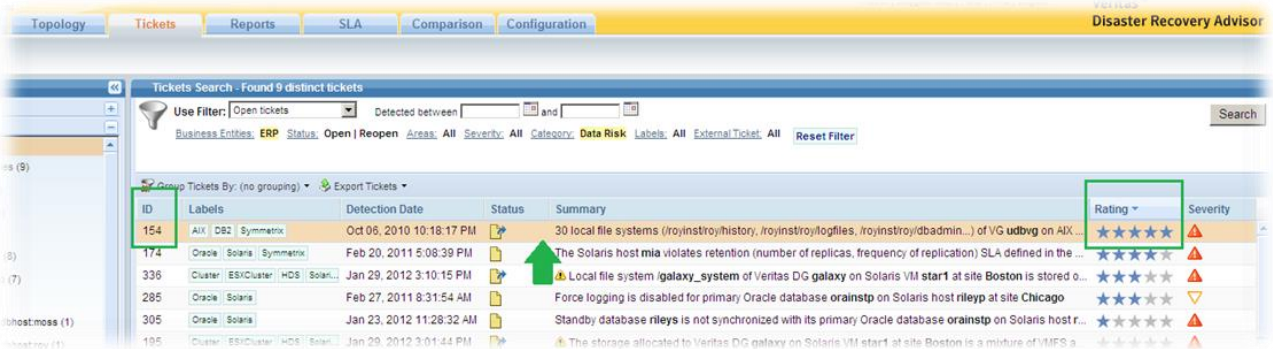

The bottom of the screen is now populated with additional information to help you fully understand the issue, evaluate the impact, review resolution options and collaborate with the right teams – all through a single, intuitive interface.
You can adjust the size of each display area by placing the cursor on the gray separator line. After the cursor changes its shape, left click and drag the separator line in any direction (see ❶ in the screenshot below).
The ticket itself contains several tabs of information (see ❷ below). The **Description** tab captures all the information you need to understand the technical issue and remediate the problem. Other tabs help you understand the specific impact of the ticket, suggest ways to resolve the issue, describe the history of the issue (has this ticket been open and closed in the past?), include space to add notes on the ticket, and list symptoms that lead DRA to discover the issue.

Finally, there are several action buttons – some of which we'll explore in this lab session. The first one is the "**Open Topology**" button (❸). This generates an interactive topology view of the ticket which is displayed above the information area. Press it now to explore the topology.

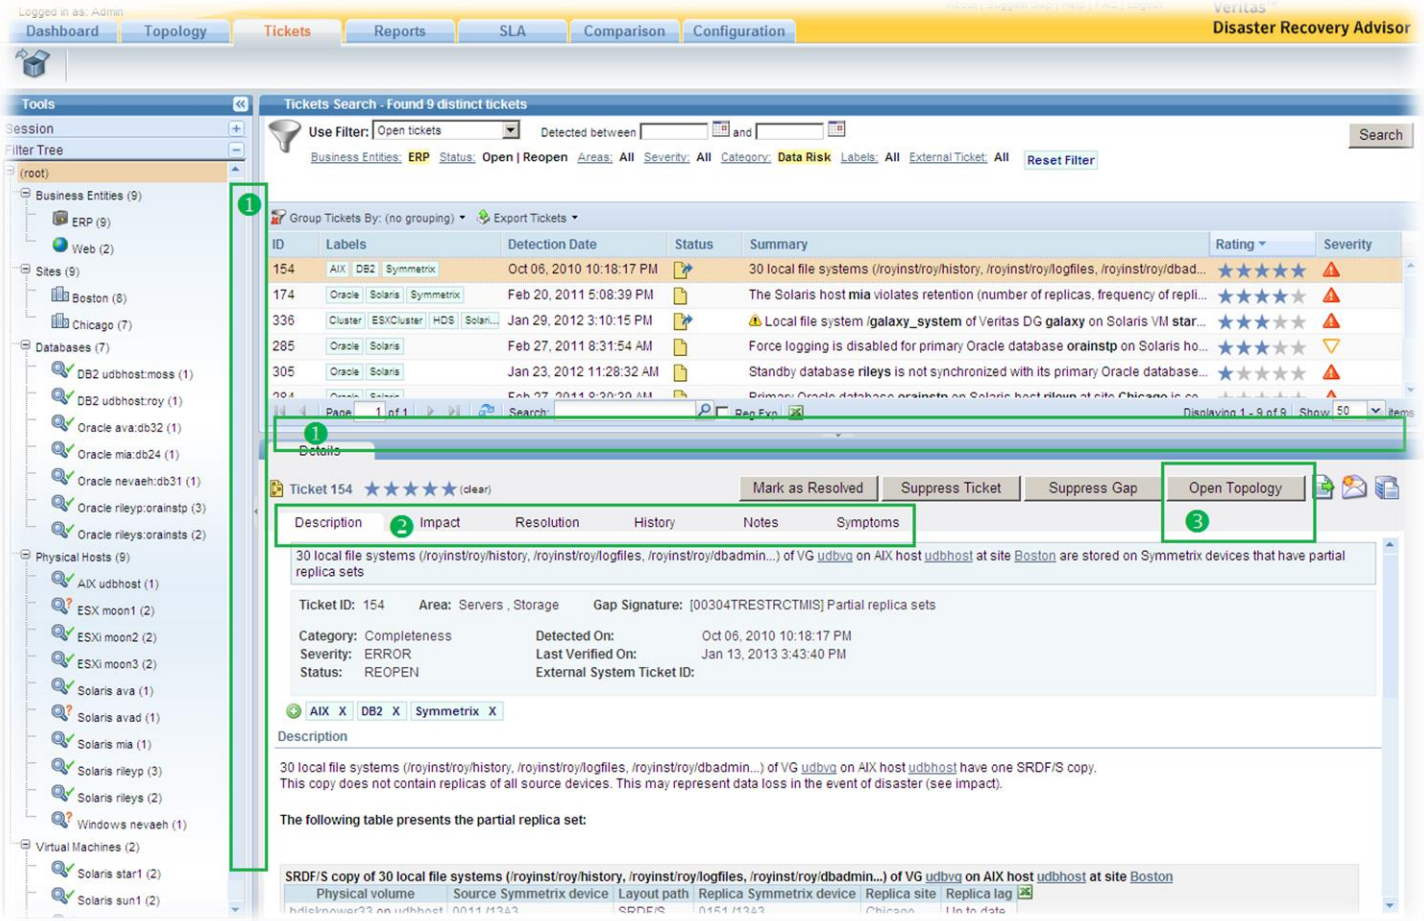

You could easily manipulate the topology that appears to reveal more information and explore additional aspects of the situation described in the ticket.
The initial view attempts to capture the essence of the problem. In this case, a UNIX volume group (VG) is stored on 12 Symmetrix devices, only 11 of which are being replicated. The problem this ticket reveals is data loss in the event of a disaster. The copy of the VG is incomplete and as data is striped on volumes – it would be lost completely.

Whenever one or more of the items can be singled out as causing a problem, they are highlighted in red. In this case, it's the un-replicated Symmetrix device 24D7.

Let's reveal more detail of this ticket:
Click on the box containing the 11 source Symmetrix devices (❶ below), press and hold the shift key, and now click the 11 target devices (❷). If done correctly, you'll see the two groups highlighted in tan. If you've selected any other item

by mistake, you can shift-click it again to unselect. Now click on the **Expand Selected Entities** tool (❸) to reveal more information.
The layout tools (❹) are useful when you realign complex topologies, and you can use the **Zoom** tools (❺) to adjust size. A wheeled mouse can also be utilized for zooming.
Panning can be performed by right-clicking, holding, and moving.
When you click any element in the topology, the ticket moves to the background and the item tab is revealed to show information about the selected item. Try clicking on one of the storage volumes, the host and even an SRDF link to see its configuration automatically display under the **Item** tab.

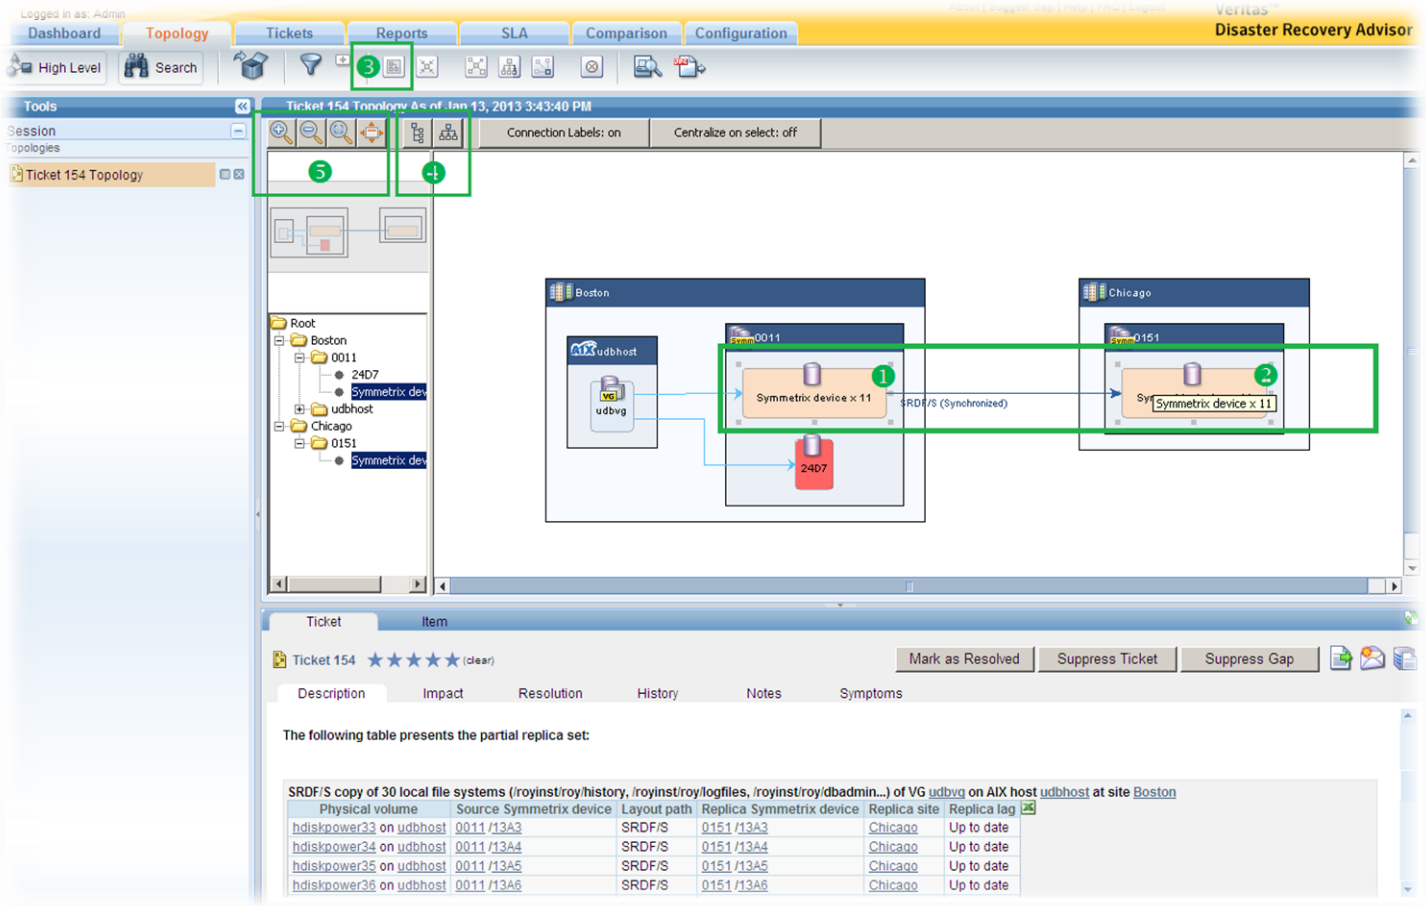

You can always bring back the ticket tab to the front by clicking on the **Ticket** label (❶ below). Let's do it now to explore one more option for interaction between the topology and the ticket panels. Use the scroll tab (❷) to move down the ticket **Description** tab until you see the databases affected by the ticket. Click on the **mossdb** instance (❸) to have it immediately added to the topology view.

Once you've reviewed the situation, you can add notes, suppress the ticket (forever or for a definite time period), and collaborate with your peers to resolve the problem. The **Send ticket to** button (❹) can be used to send the ticket to selected recipients, and the Send to external ticket system button (❺) can used to immediately push the ticket information to an external system of your choice. Of course, these two tasks could be fully automated for all tickets or just for tickets that meet certain criteria (optional Lab exercises 6-7 will walk you through these tasks).

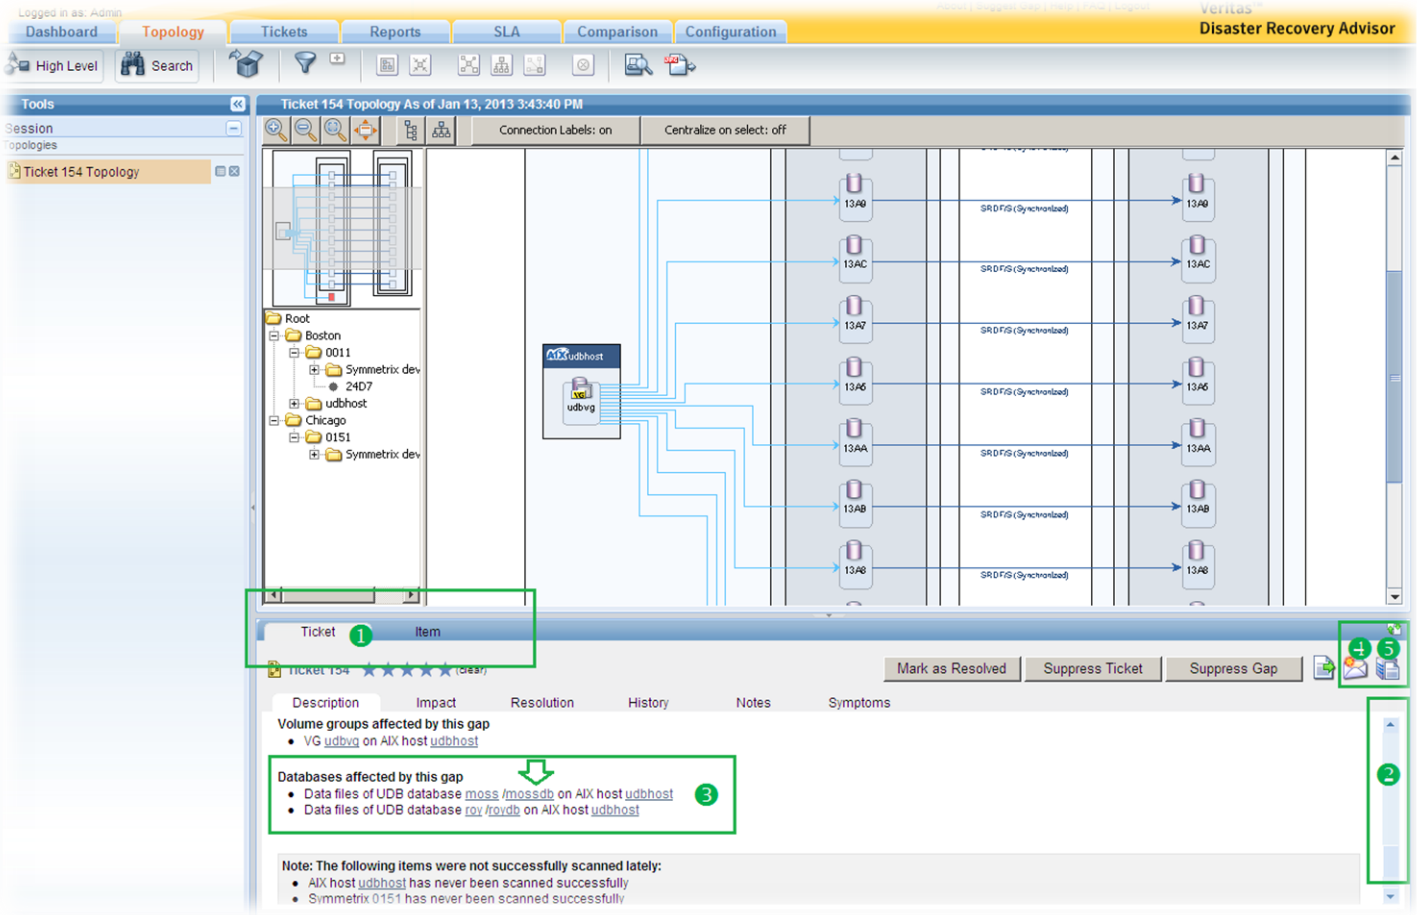

Lab Exercise 2

Topic – Configuring Comparison Analysis	Duration ≈ 10 minutes

The scenario for this next lab exercise is the following:

- A fail-over in a mission-critical Sun-cluster occurred
- The application is behaving strangely, with intermittent UI failures

Let's run a comparison of the two nodes to highlight differences and see if we can pinpoint what is causing the problem.

Select the **Comparison** tab at the top of the DRA user interface (❶ below).

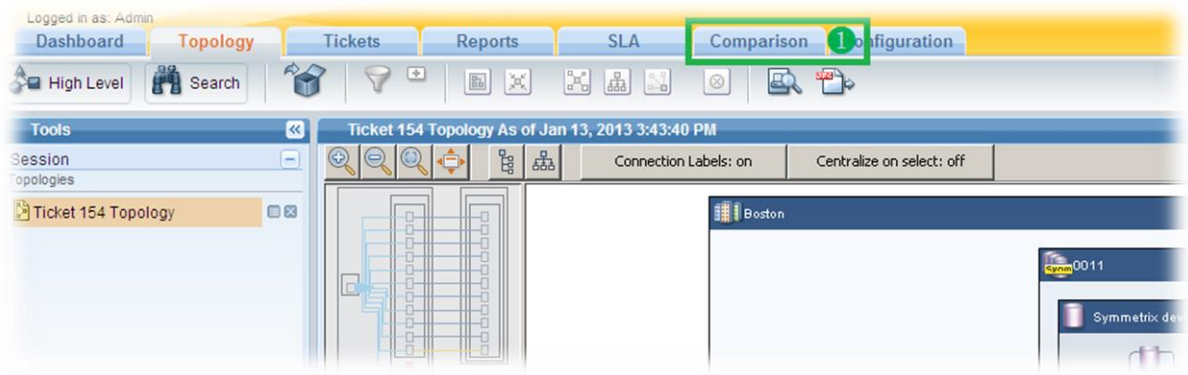

You are directed to the comparison tab within DRA. From here we can build a worksheet to compare different hosts or clusters and even groupings of each. Let's begin by clicking the "**Create new Worksheet**" tab at the bottom of the page.

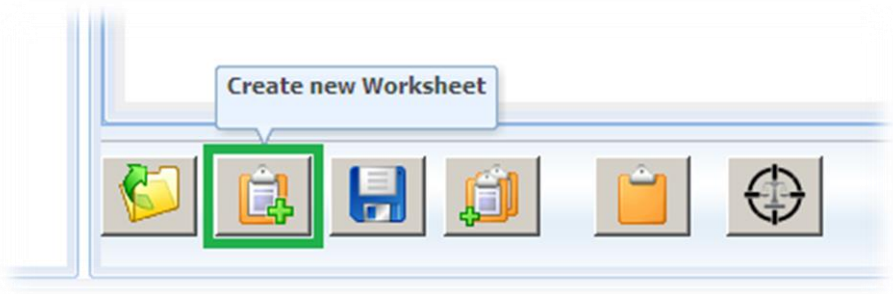

As you can now see, you have begun to create a comparison worksheet that will track each comparison you run. You can name the worksheet (❶) and also specify who should be able to access and/or manipulate the contents (❷).

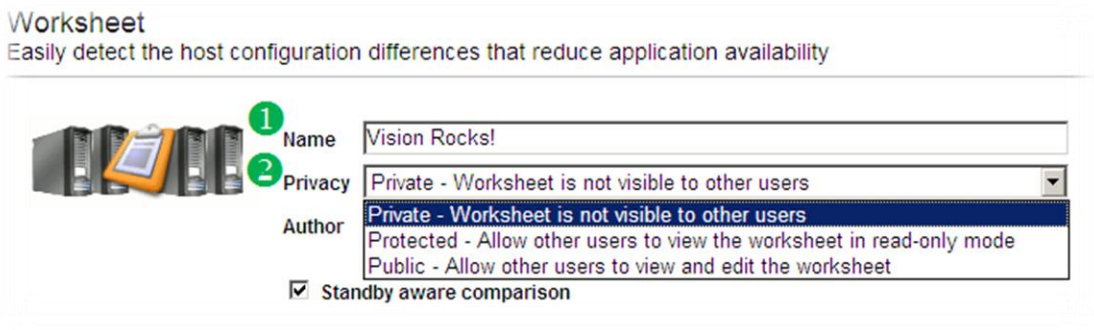

Various options are now available to configure your comparison worksheet. The first option below (❶) is used to create a custom comparison grouping.

The second option (❷) allows you to compare your High Availability clusters. We will select option two momentarily.

Finally, option three (❸) allows you to compare clusters with their Disaster Recovery counterparts.

Let's dive into the second option. Click on **option** ❷ (above) to bring up the comparison configuration module.

Select where it says (none) in the clusters. In the next dialog box, select all types of clusters by clicking the box next to **Type** (❷). Then select **OK** (❸). Finally select **Create** to populate.

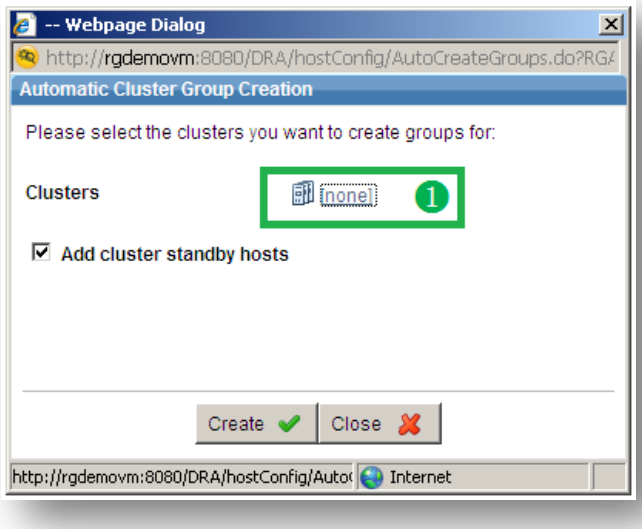

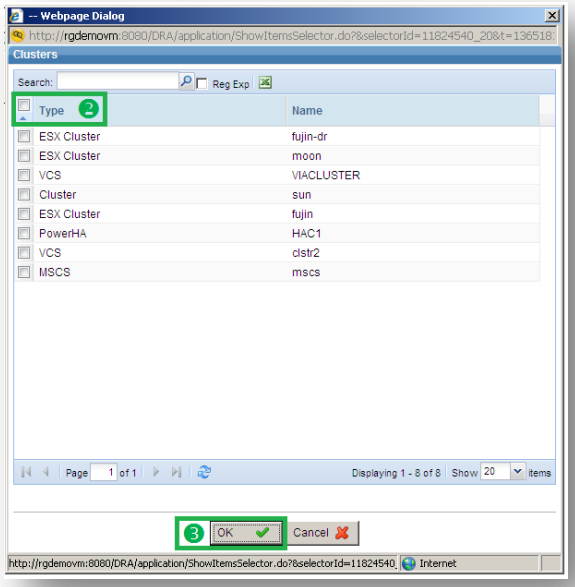

Click on the **Sun** cluster (❶ below) and then click on the **Sun** tab. Now select **Compare!** (❷) to run the comparison.

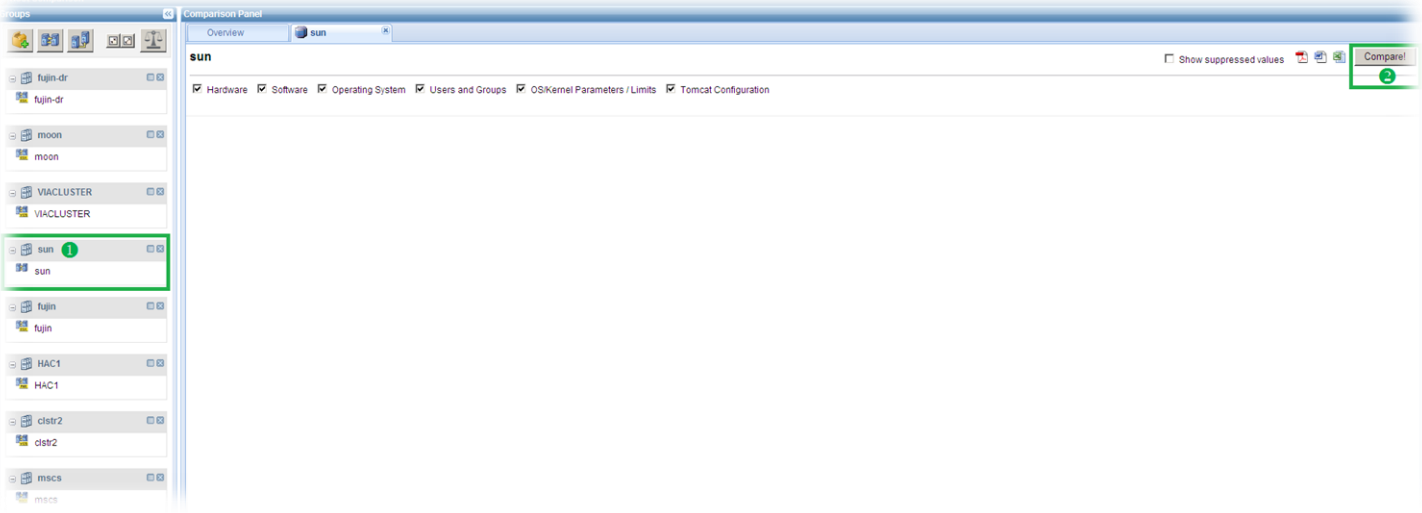

DRA is preconfigured to look for specific configurations to compare (❶ below). Those configurations include Hardware, Software, Operating Systems, Users and Groups, and OS/Kernel Parameters/Limits. You also have the ability to add specific comparisons to the group. In this demo environment, we have added tomcat configuration.

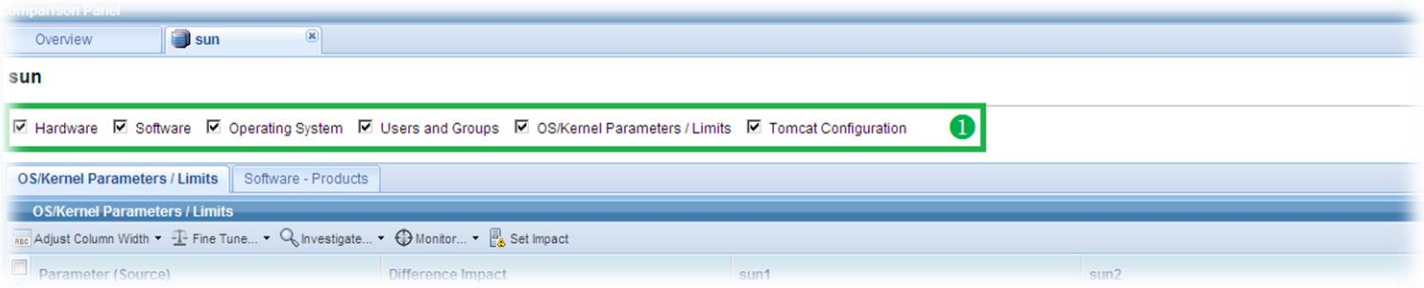

The tabs in this area (❶ below) display the differences that have been identified in the cluster by DRA. You can see that differences have been found in two areas, **OS/Kernel Parameters/Limits** and **Software – Products**.

The first tab shows some interesting kernel tuning differences. The real issue however, is in the **Software – Products** tab. Click on the **Software – Products** tab to expose the differences. Choose the **Adjust Column Width** option and then select **Content and Heading**. DRA has identified, among other differences, that the Java Extension Libraries are different within the cluster. This has been identified in the past as a common code for the application issues we now experience.

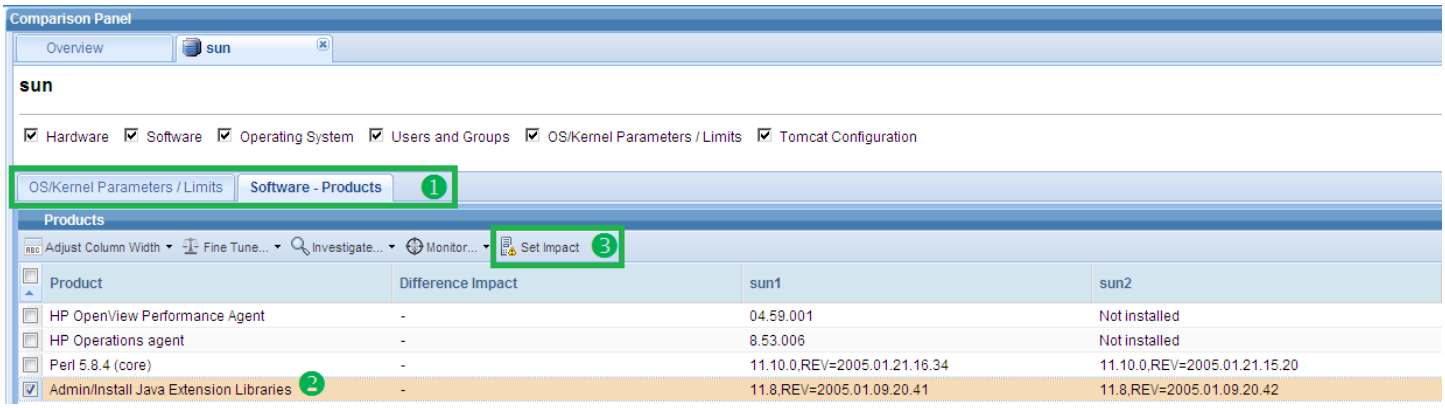

Let's go one step further and configure DRA to proactively identify and alert you of similar misconfigurations.

To do this, first document the impact of the identified difference:

Check the box next to the **Admin/install Java Extension Libraries** and select **Set Impact** (❸ above).

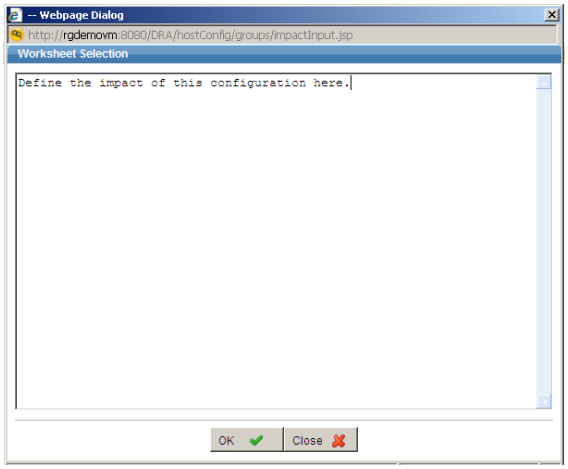

In the dialog box that appears, you can define the impact this issue has created. In this case, the misconfiguration in installations of Java Extension Libraries has brought down one of your nodes in the cluster. There is room to add as much information as necessary.

Finally, check the box next to the **Admin/install Java Extension Libraries**, click the **Monitor** option (❶ below) and select **Monitor changes in tickets module.** If preformed correctly, a small target will appear next to the selected misconfiguration. Now DRA will provide you with an open ticket any time this issue is identified within your clusters.

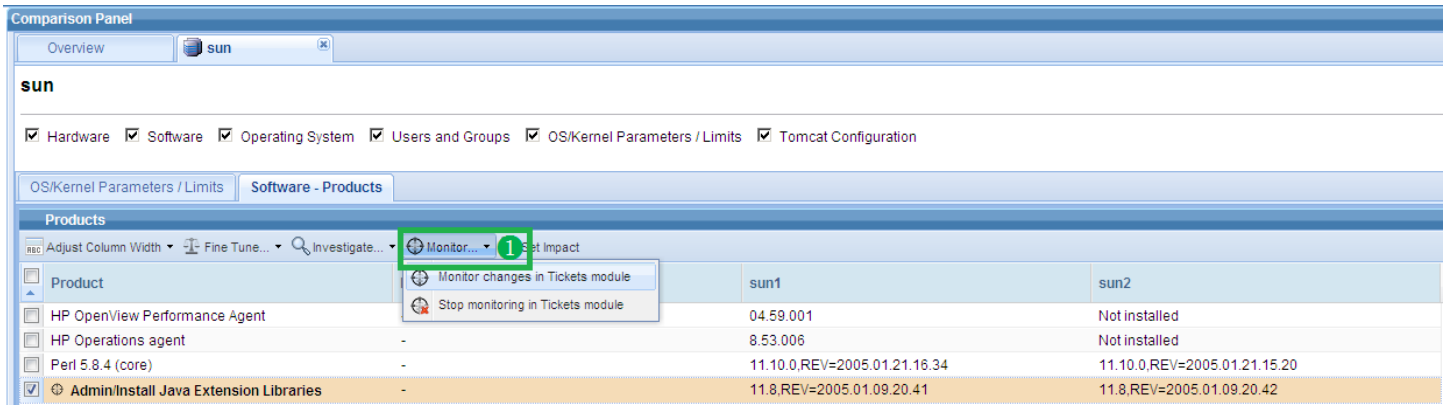

Lab Exercise 3

Topic – Proactive Risk Detection in the Private Cloud	Duration ≈ 10 minutes

The scenario for this next lab exercise is the following:

- Various components of the ERP business service are being gradually virtualized
- To make sure this does not introduce unnecessary (or unknown) risks, the application owner must regularly review the status of the virtualized components

Note: You can configure this process to export the reports automatically (see automatic exporting). However, we will walk through the process of manual reporting in this next exercise.

Click on the **Tickets** tab (❶). Click **Reset Filter**.

Use the top search bar to narrow the displayed tickets only to those relevant to the virtualized infrastructure. Click on the Labels link (❷).

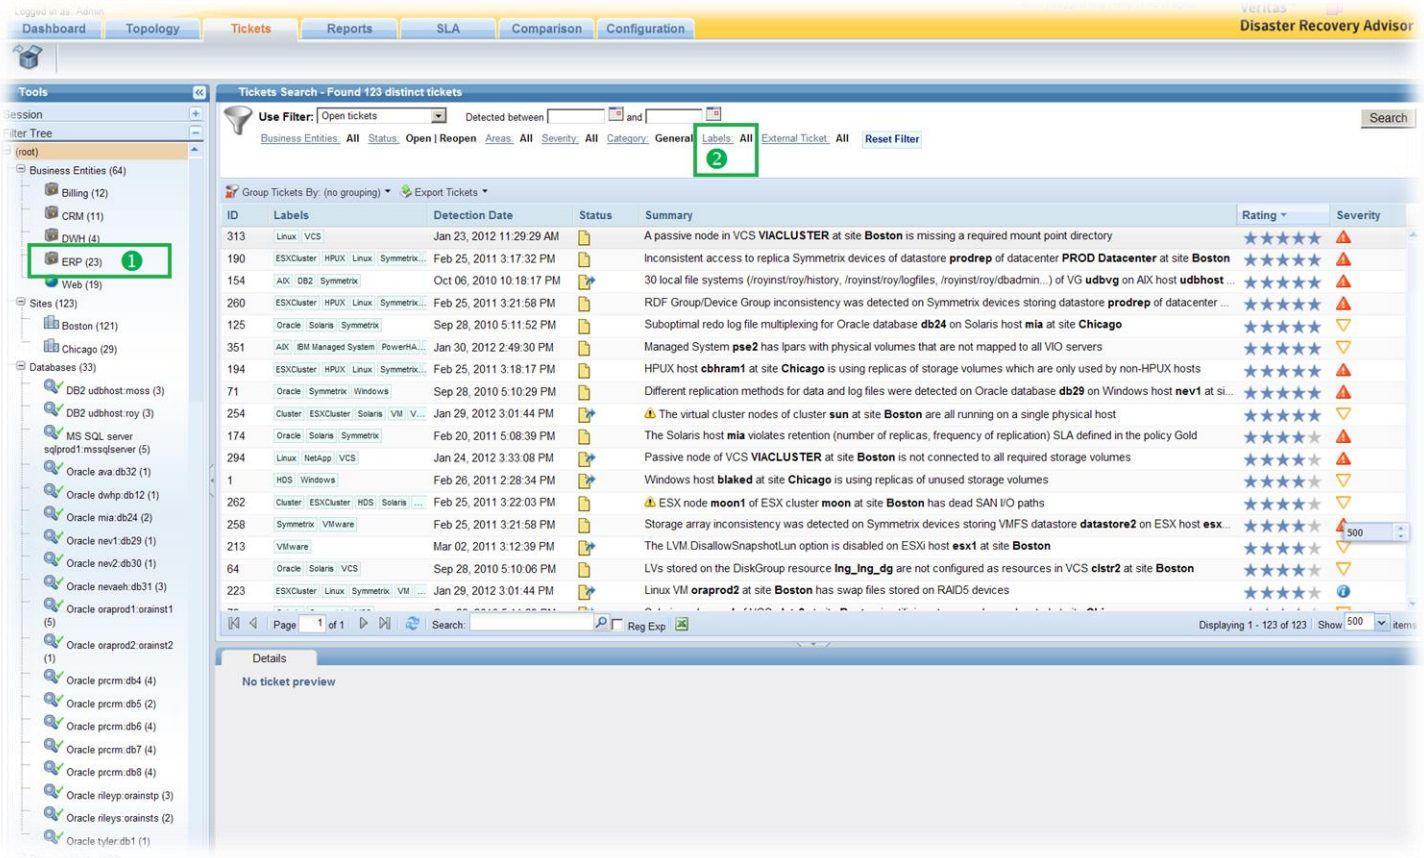

An advanced search dialog box will appear, with the **Areas and Labels** tab exposed (other tabs allow you to refine other search options). Select **Filter by Ticket Label** (❶), scroll (❷) to the bottom and select **VMware** (❸). Click **OK** to complete your selection.

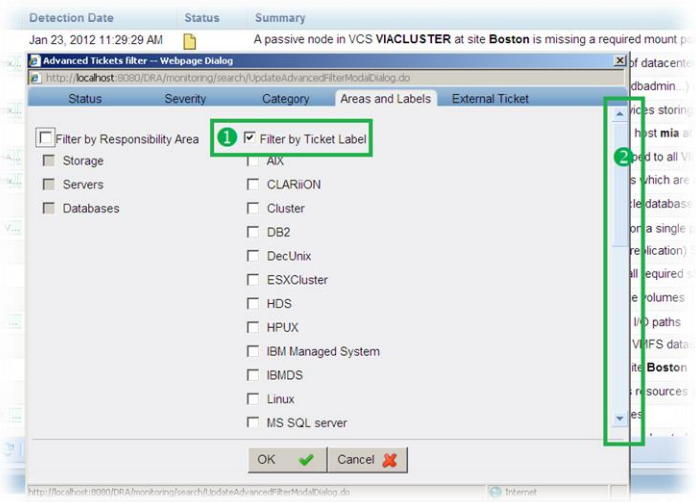

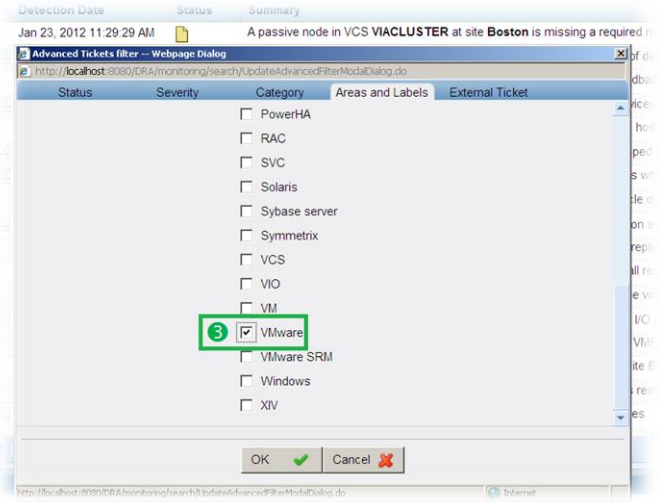

Click **Search** (❹ below) to load the new search criteria. Filter the viewed tickets to show only **ERP** related tickets by clicking on the corresponding business service on the left **Filter Tree**. You now have a clear view into all the ERP-related VMware virtualization risks.

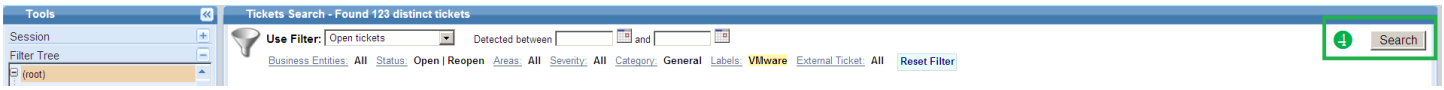

We can see that there are six risks identified. We will explore the two with the highest rating. Feel free to experiment with the rest. Click anywhere on the first ticket, **ID 254** (❶), then open the topology (❷).

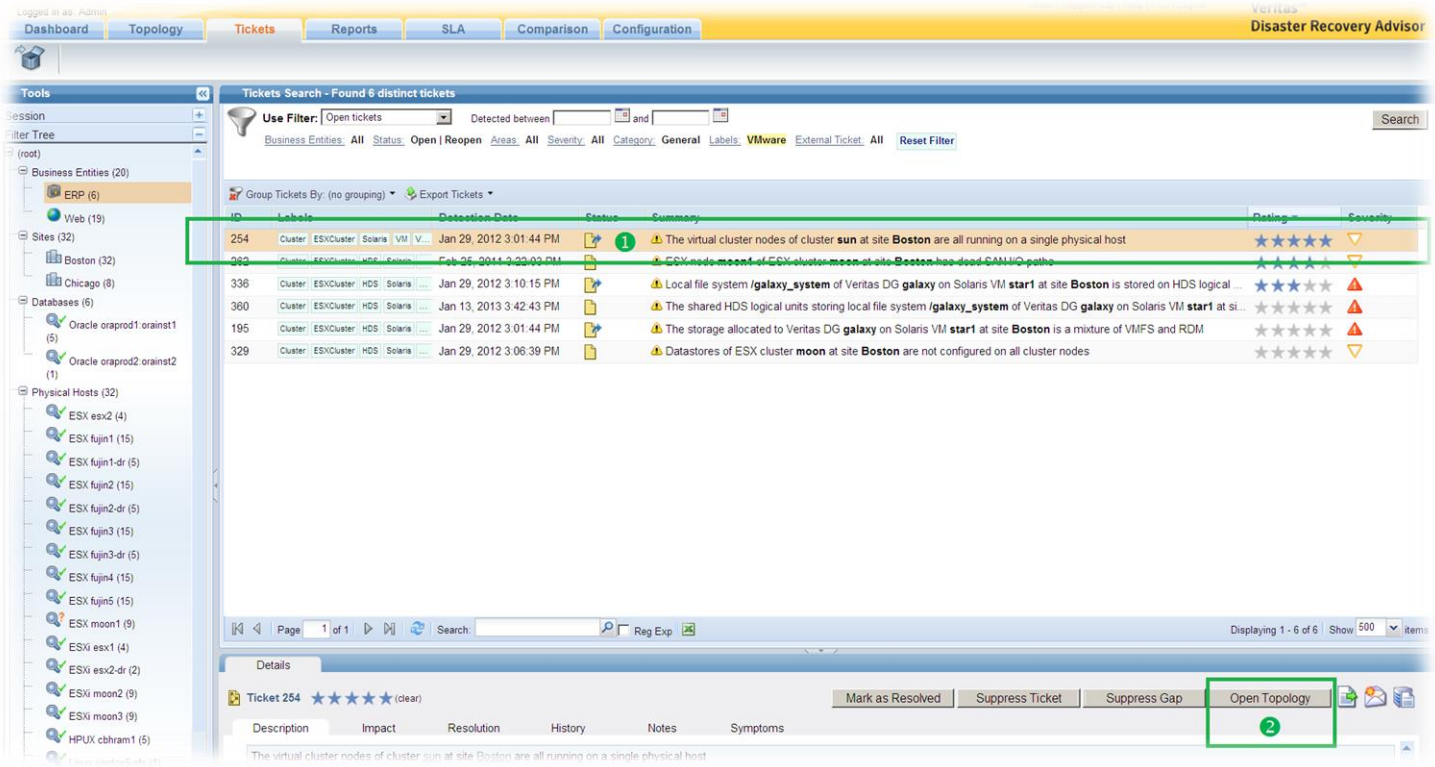

Try one of the layout buttons (❶ below) and zoom controls (❷) – or your mouse wheel to change the aspect and zoom.

The ticket reveals an interesting situation: two virtual machines (❺) – likely physical servers that have been migrated to the private cloud – are configured in a cluster. Both VMs are currently running on the same ESX node (❹). This, obviously, introduces a single point of failure that puts the service at risk.

This situation can develop over time, as VMware Dynamic Resource Scheduling (DRS) could periodically move certain VMs from one ESX node to another to better load balance resource utilization, or as a result of a human error.

If you're interested to get advice on possible ways of avoiding this situation in the future, consult the ticket resolution tab.

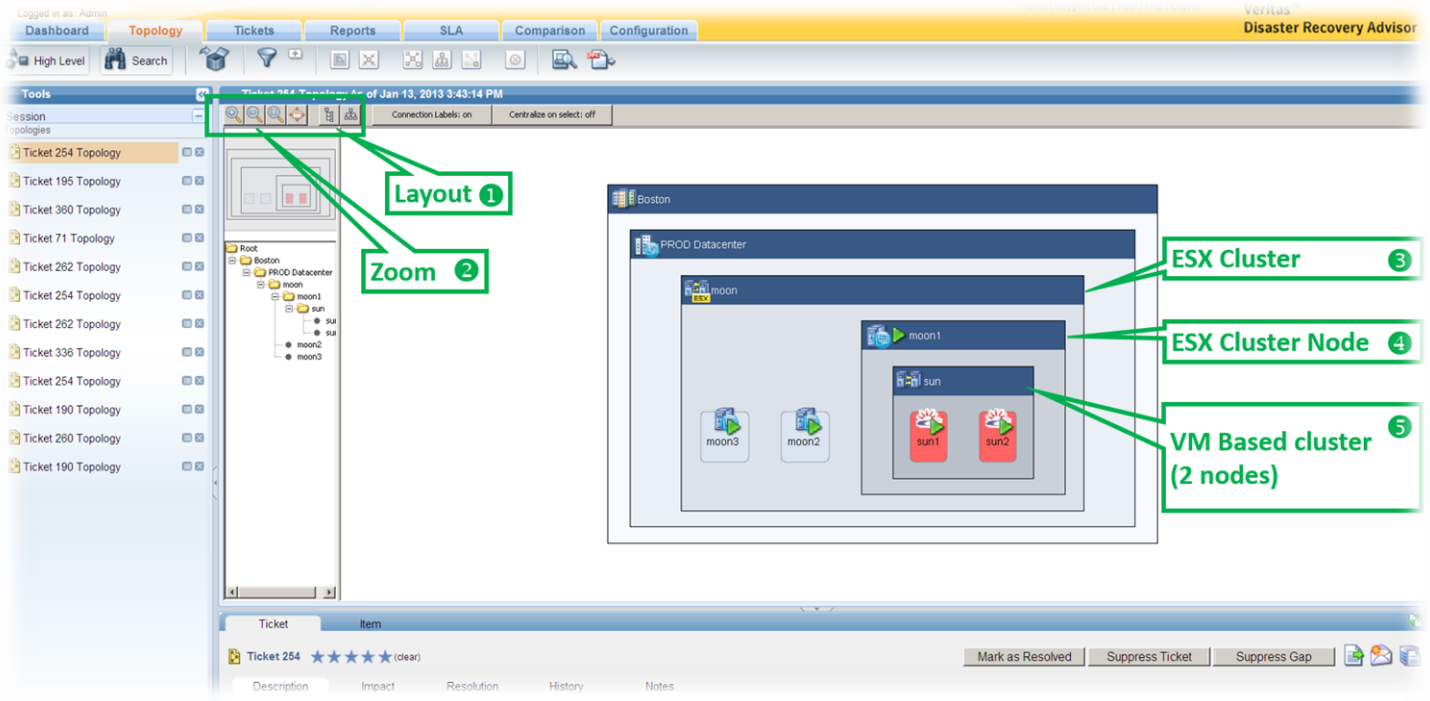

You can return to viewing the other VMware admintickets on the ERP environment simply by clicking on the top **Tickets** tab (1 below).

We encourage you to review one or more of the other tickets (details tab, topology, etc.) – for example:

- **Ticket ID 262** – Dead SAN I/O path brings a cluster to single point of failure
- **Ticket ID 195** – VM allocated with storage from different storage tiers (RDM and VMDK), both used in a striped Volume Group (could lead to serious data corruption!)
- **Ticket ID 329** – ESX cluster configured with Datastore that are not mapped to all ESX nodes. This could break the VMware HA mechanism.

Lab Exercise 4

Topic – Detect Recovery Point Objective (RPO) SLA violations	Duration ≈ 10 minutes

The scenario for this next lab exercise is the following:

- Tier-1 applications must meet an RPO of 5 minutes or less. Other tiers have different RPO requirements (out of the scope of this exercise).
- Prior to owning DRA, there was no simple way to guarantee that the required RPO SLA is met at all times.
- As a result, this limitation regretfully lead to unexpected loss of data, prolonged recovery and failed DR exercises.
- With DRA it is possible to:
  - Define what the required RPO is and automatically receive an alert when it is violated
  - Proactively review the actual RPO of an application and its underlying infrastructure

- Other than RPO, DRA SLA policies enable defining various HA/DR standards that will be enforced by DRA such as I/O path redundancy, DR capacity, number of snapshots and more

View and Define SLA Policies

Select the **Configuration** tab (❶ below) and then click **SLA Definition** (❷) on the left navigation tree. The list of SLA policies defined by the user is now presented. Click on the **Gold** SLA policy (❸).

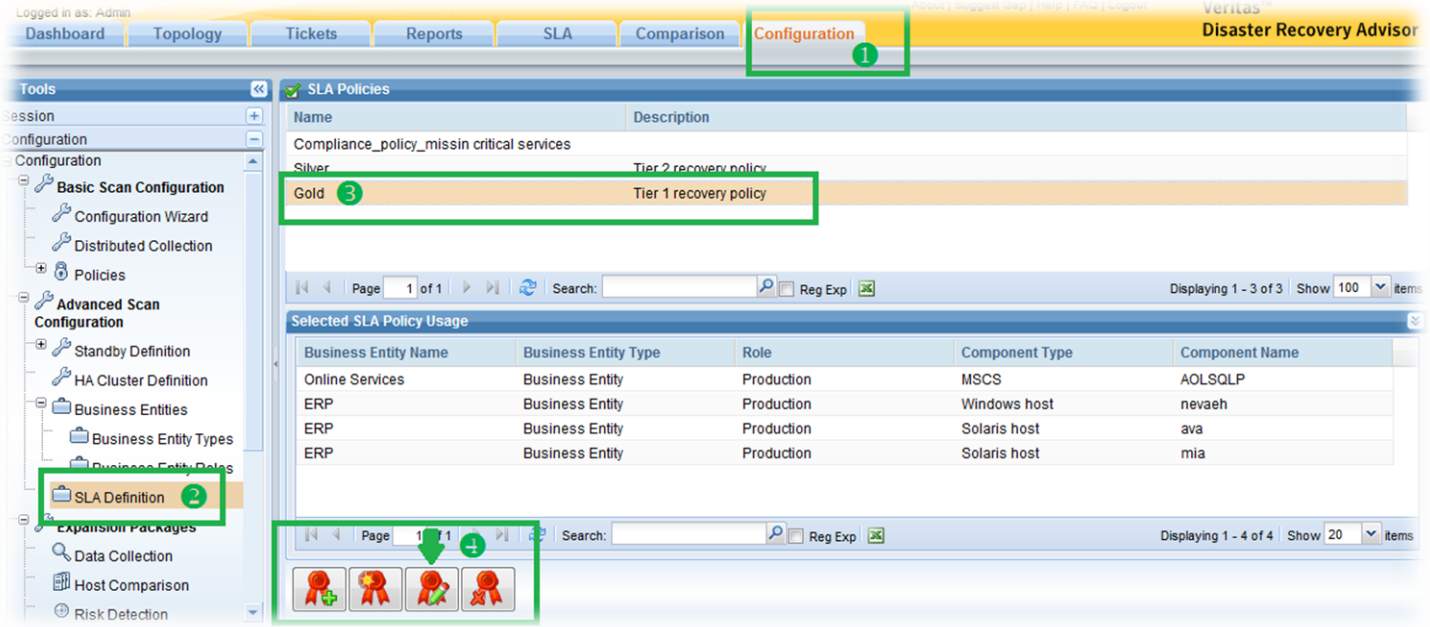

The bottom pane now presents which business entity components the user has associated with the selected **Gold** SLA Policy. The Gold SLA Policy is associated with an Online Services Microsoft Cluster (AOLSQLP) and several ERP Systems. The buttons at the bottom of the screen allow users to Add, Clone, Edit and Remove SLA Policies correspondingly. To view the set of standards defined in the **Gold** SLA policy, click **Edit Policy** (❹). You may also double-click on the Gold SLA policy record.

A new window is opened presenting the SLA policy properties. Note that there are four tabs – **Data Protection**, **Availability**, **Standby/DR Capacity Plan**, and **Database Replication**. The **Data Protection** tab (screenshot below) shows that the user defined to enforce an RPO of 5 minutes at a remote site. In other words, any business entity component associated with this policy is required to have a replica in a remote site which is 5 minutes old (or less). While it is possible to define the required replica type (SRDF, HUR-ShadowImage, etc.), the user did not define it explicitly in this case. Note that the **Enforce** checkbox is selected. This means a ticket will be opened when DRA detects a violation as part of the scheduled risk analysis (typically scheduled on a daily basis).

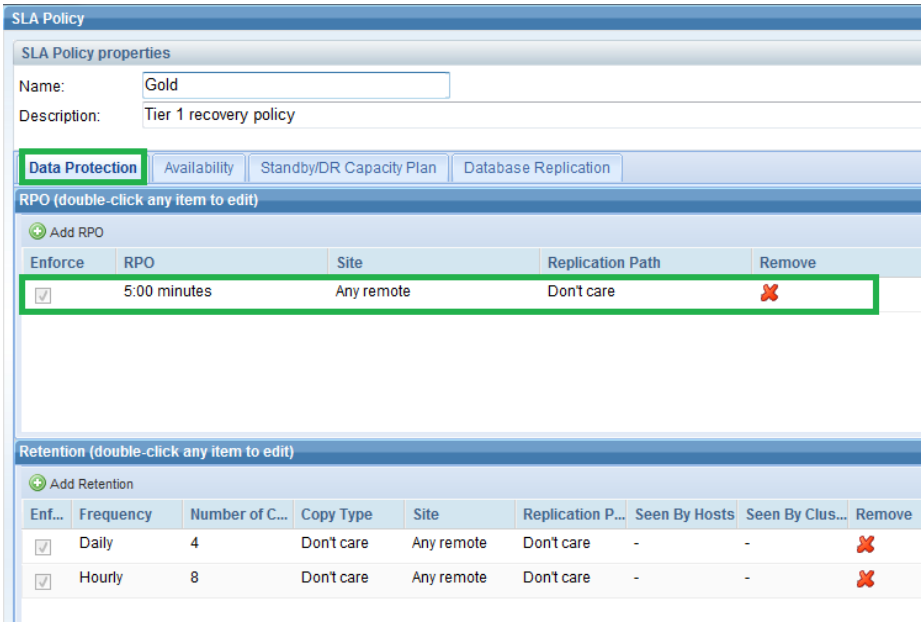

Feel free to browse the different tabs to view other standards defined by the user and enforced by DRA as well as other available standards.

Note: While not in the scope of this exercise, DRA offers a simple way to associate business entity components with SLA policies through the Business Entities page available in the Configuration tab.

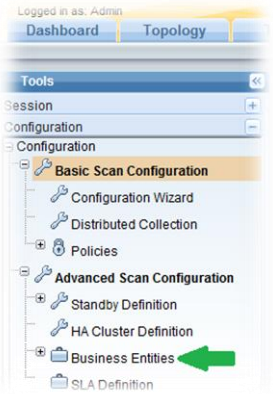

Now that we learned about SLA policies definition and association with business entity components, we can continue learning how to govern compliance with SLA policies and view open SLA risks.

Get immediate indication for SLA breach

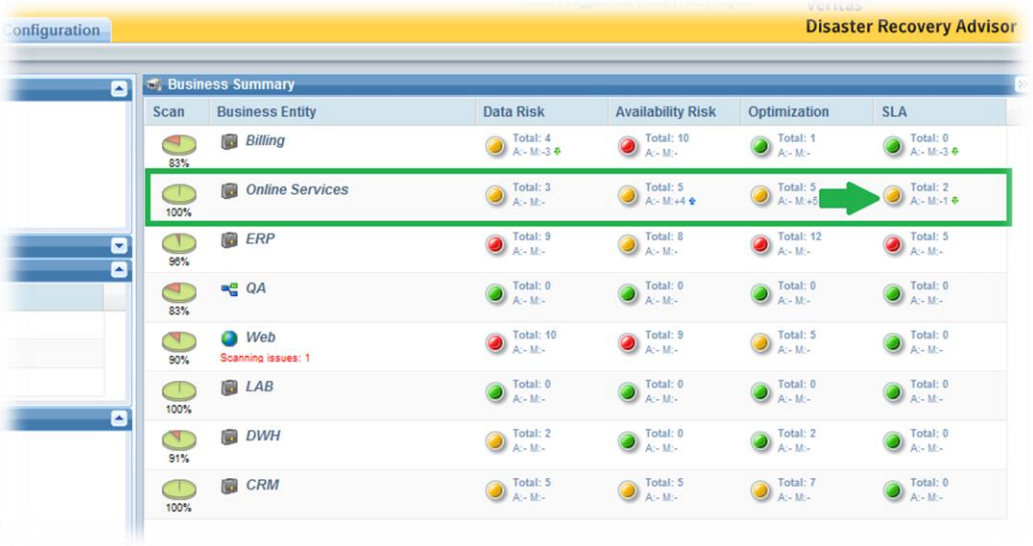

- Enter the **Dashboard** tab to View the high-level business status of the **Online Service** business Entity
- Notice the risk indication under the **SLA** column. A total of 2 risks exist.
- Next, we'll want to get better insight into the identified risks and take action.
- Click on the yellow risk indicator (●) under the **SLA** column to view more details

Faults and Risks detailed information is now presented (screenshot below), filtered by business entity (**Online Services**) and category (**SLA Breach**). By default, the tickets are grouped by **Gap Type**.

Note: if "**No Grouping**" was defined in a previous lab exercise, click on **Group Tickets by** and select **Gap Type**.

A Gap Type represents one of many risk signatures available in the DRA risk knowledgebase. Note that two different SLA violations were detected – **Minimum number of I/O paths** and **RPO SLA**. Expand the **Host violating RPO SLA** ticket group.

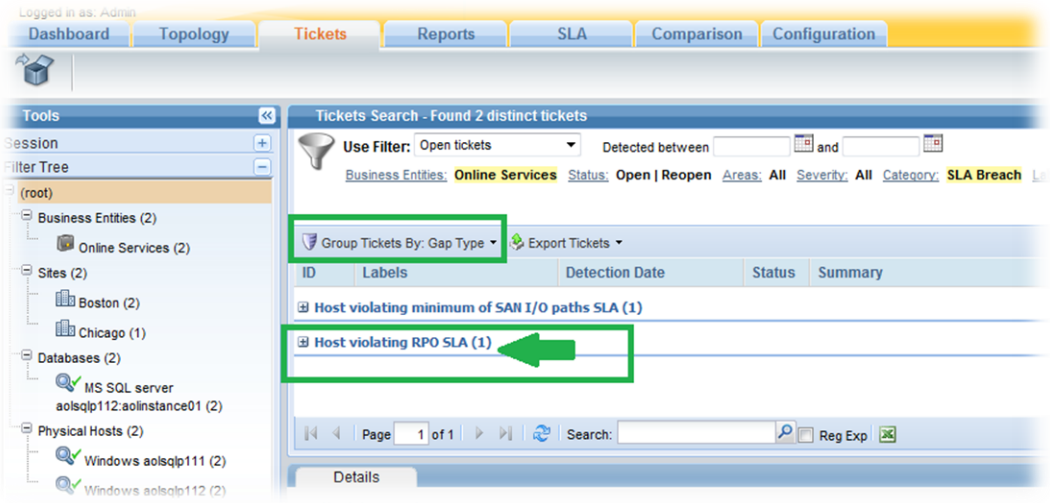

The ticket list reveals an **RPO SLA violation risk** detected for Windows cluster node aolsqlp112. More details will be visible immediately upon clicking any of the tickets in the table. Click on the ticket summary for ticket **ID 556** (❶ below).

The bottom part of the screen is now populated with additional information to help you fully understand the issue, evaluate impact, review resolution options and collaborate with the right teams – all through a single, intuitive interface. In the Ticket **Description** tab, click on the **Expand/Collapse Table** symbol (❷) to unfold the list of affected components.

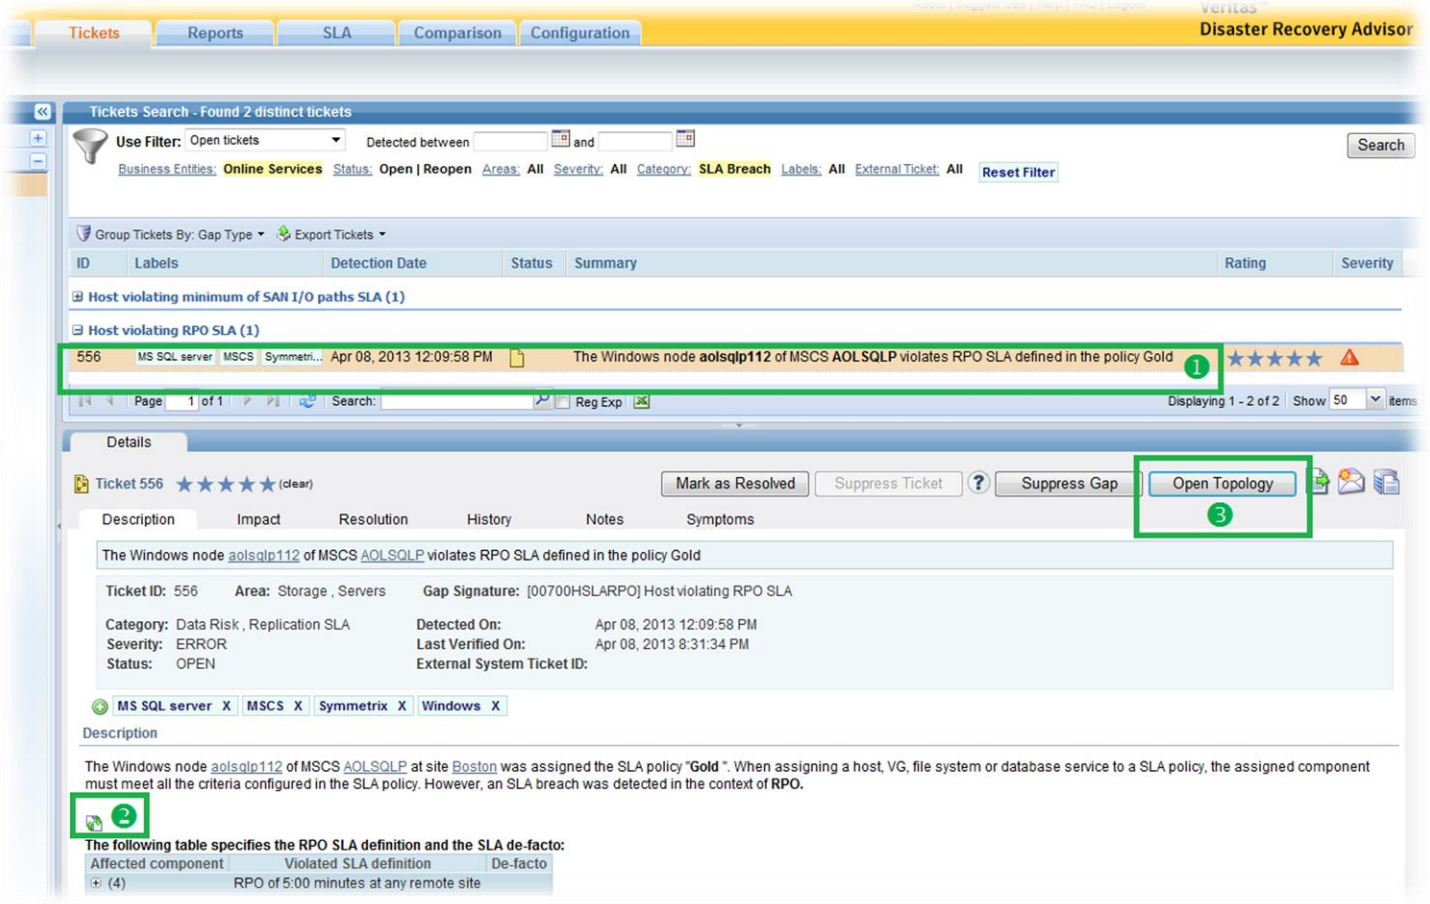

The expanded table identifies the affected components, the required RPO and the actual (de-facto) RPO. In this case, 4 disk drives are expected to have a remote replica with RPO of 5 minutes or less whereas de-facto DRA found that 3 disk drives have a remote SRDF/A replica with RPO of above **17 minutes** and one disk drive is **not replicated at all**. Continue to explore the ticket by clicking **Open Topology** (❸).
The drawing reveals that the state of the SRDF/A replication is Consistent – and yet the target R2 device falls behind the source R1 device and fails to meet the required 5 minutes RPO SLA. Additional details regarding the SRDF/A replication are available by clicking on the SRDF/A connection in the topology.

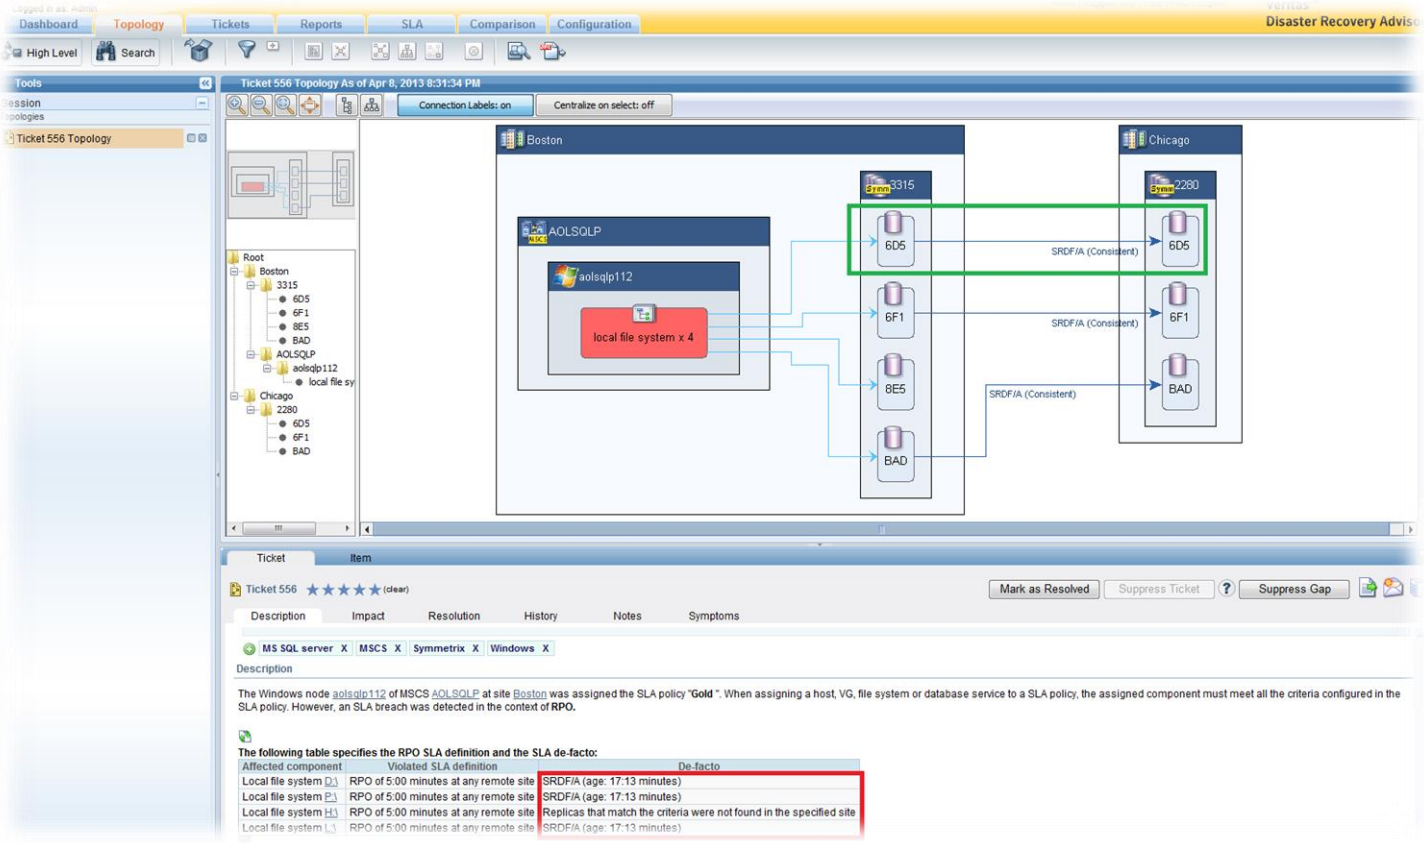

Last - feel free to return to the list of SLA tickets and explore a different violation – **Minimum number of I/O paths** (ID 636). The ticket presents a single point of failure which results in higher risk of downtime as well as degraded performance.

Lab Exercise 5

Topic – Cross Domain Collaboration	Duration ≈ 10 minutes

The scenario for this next lab exercise is the following:

- We are experiencing poor performance in one of our business services
- There were issues in our latest DR test relating to the database (db29) of this service
- The storage team claims the database is located on high-performance devices
- The DBA claims it has been configured by the book
- The server administrator has not found any configuration issues

Let's see how DRA can provide collaboration between these various teams:

Begin by navigating back to the **Tickets** tab (❶) within DRA.

Reset the filter (❶) and click **Search** (❷).

On the left filter tree, locate **Databases** (❸ below), find and select **Oracle nev1:db29** (❹ below).

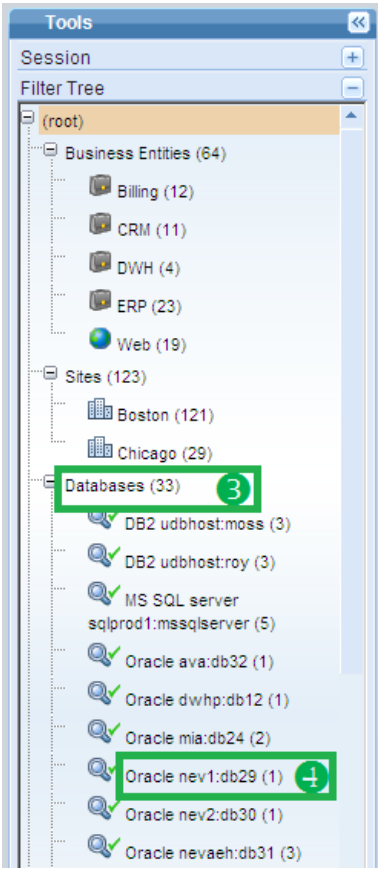

DRA has identified an issue relating to this database. Let's explore ticket **ID 71**. Click anywhere on the ticket to display the information of the ticket below.

Note: If you have not removed the grouping, simply click on the grouping title to display the ticket.

Explore the **Description** (❶ below), **Impact** (❷ below), and **Resolution** (❸ below) of the ticket to get a feeling for the issue DRA has identified. Finally, open the **Topology** (❹ below) to view the issue.

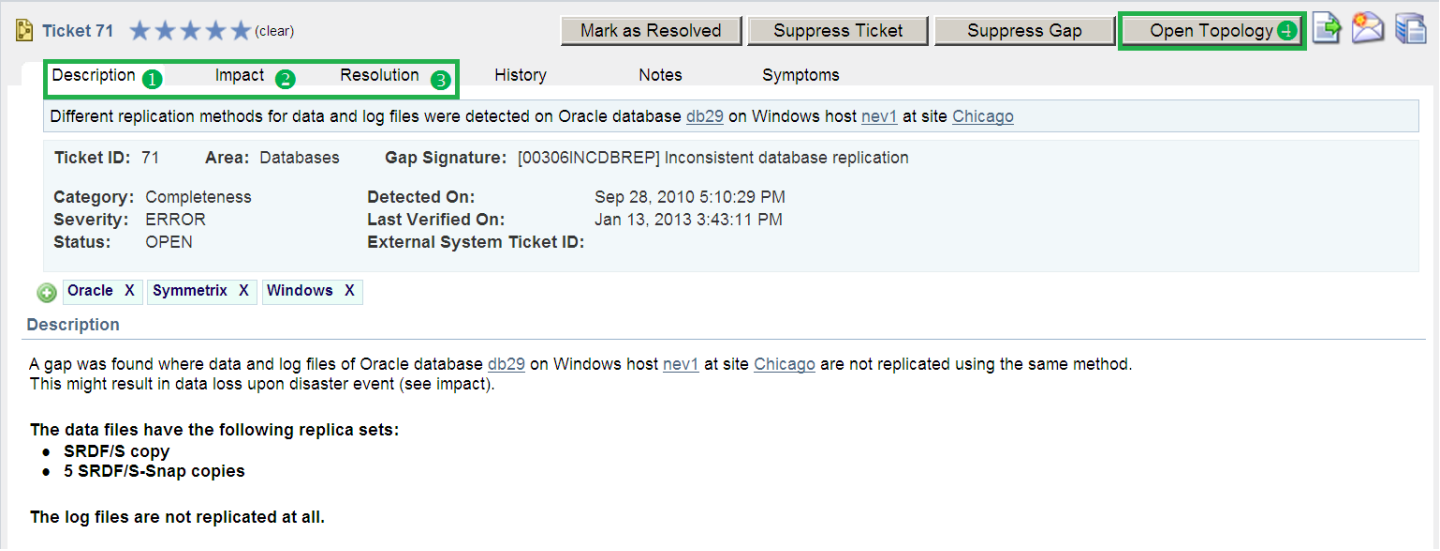

As we can now view from the topology, we have Oracle files that are currently routed to 3 different local file systems. Let's expand the topology to get a better view of how these files are configured. Select the **local file systems** (❶), then click **expand selected entities** (❷).

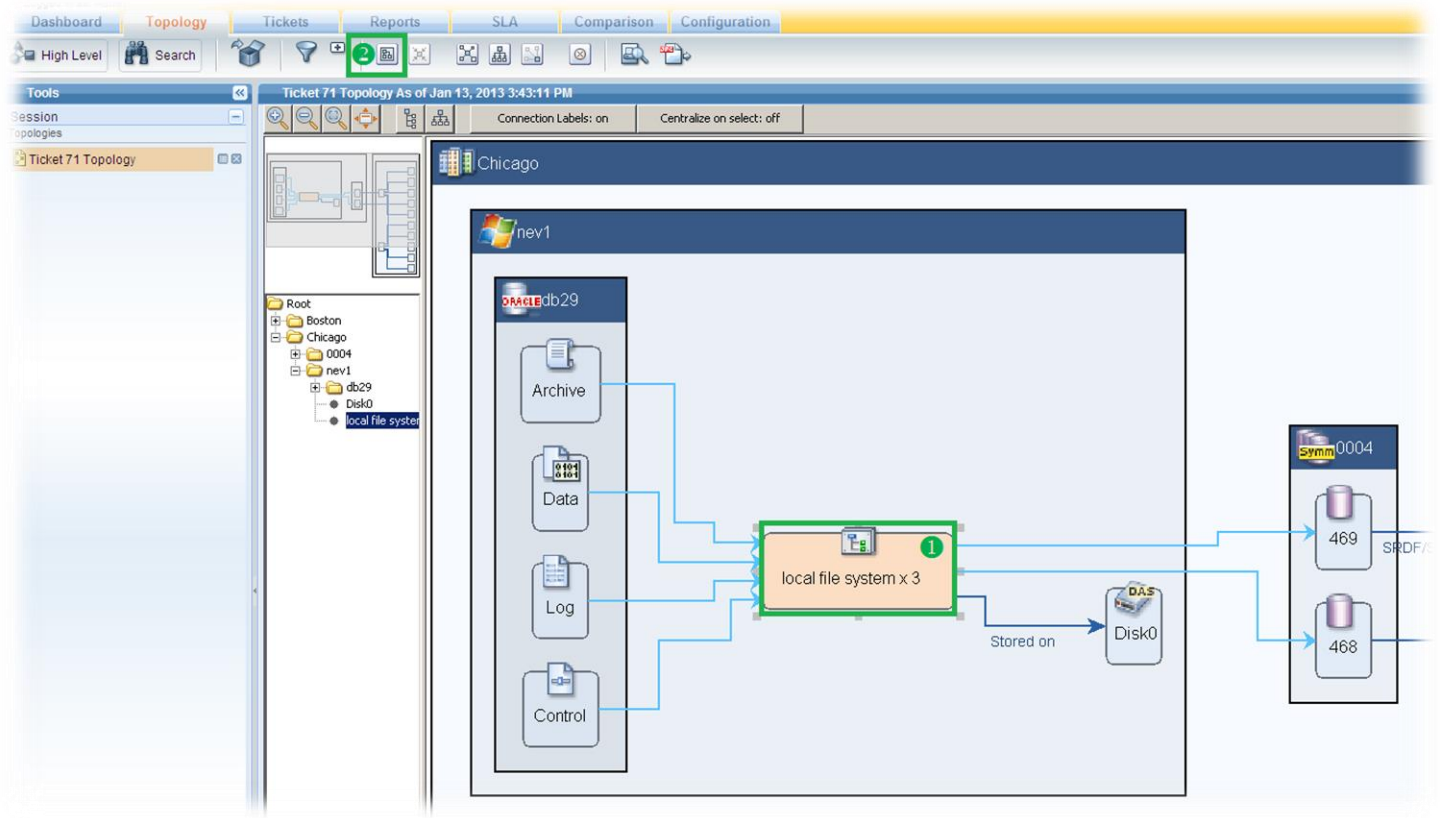

Now we can see that the Data files are located on a separate file system than the Log, Control, and Archive files. This configuration meets Oracle best practice recommendations. The issue that DRA has identified is in replication. The Data

files have been appropriately mapped and configured for replication, however, the Control, Log, and Archive files are currently mapped to poorly preforming local storage (❶ below) that is not configured for replication. In the event of a failover, the control files will not be present.

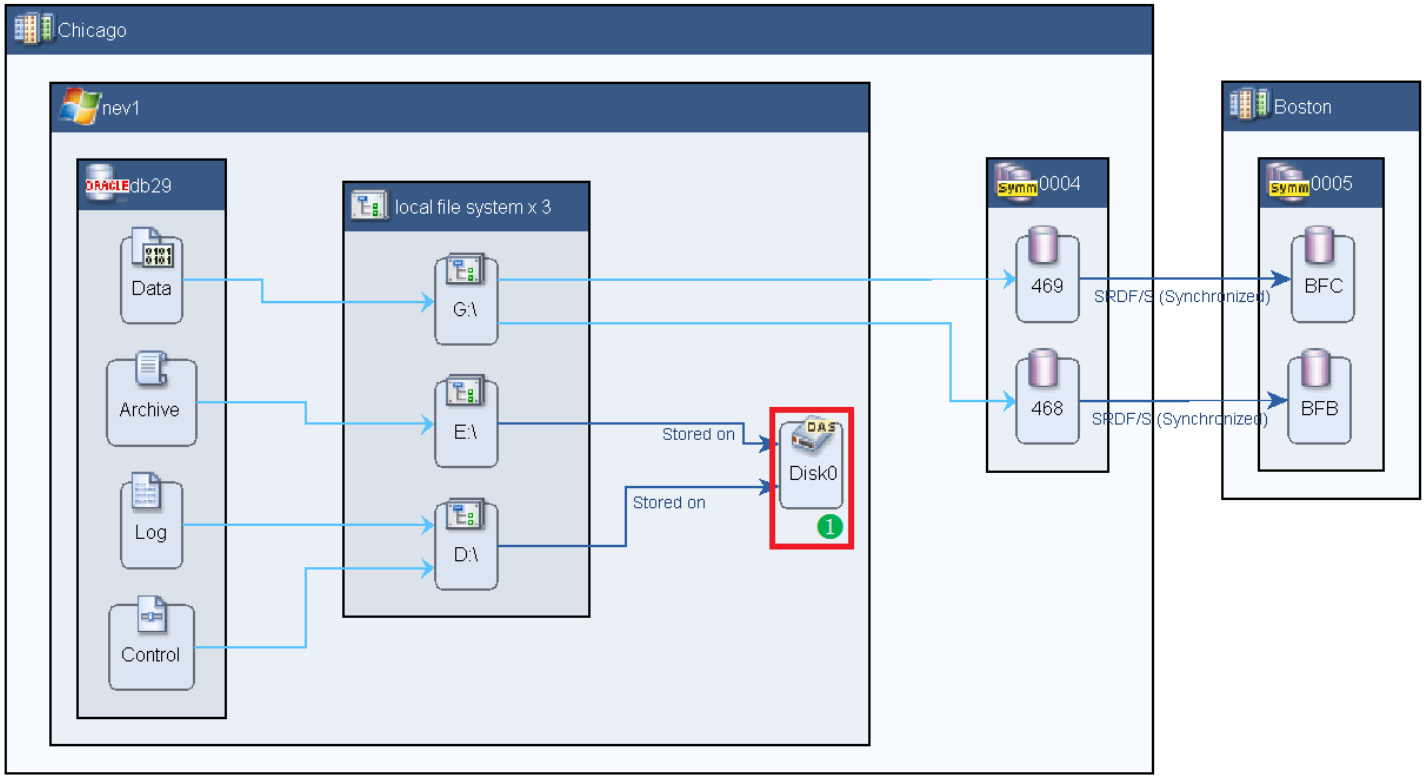

Lab Exercise 6

Topic – Create Customized, Automated Email Notifications	Duration ≈ 10 minutes

This exercise will walk you through the steps to set up automatic export of all or specific tickets:

Select the **Configuration** tab (❶).

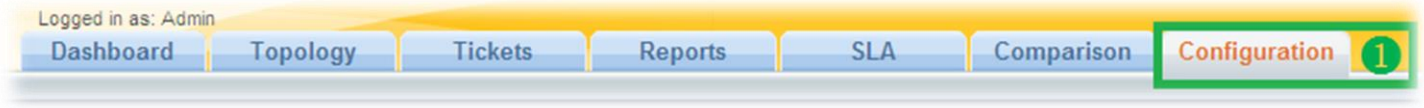

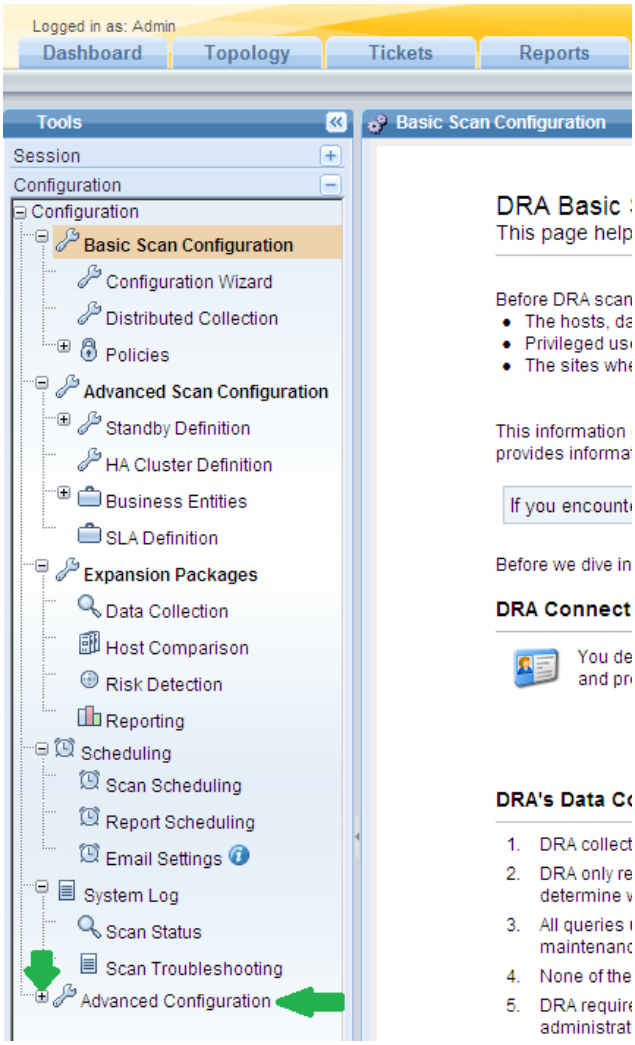

- Locate the **Advanced Configuration** option at the bottom of the left filter tree.
- Click on the + to expand the additional options
- Select **External Ticketing** from the list of options that appears.

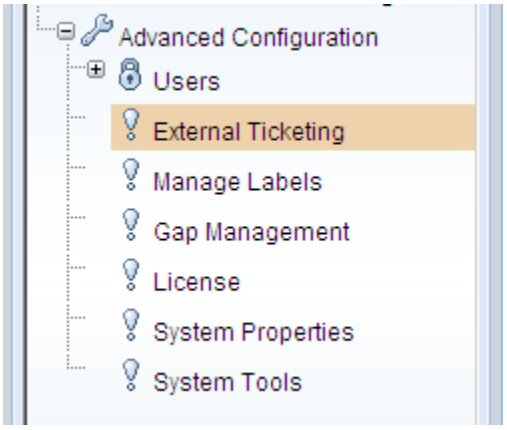

Connection mode specifies whether the export should be sent manually or automatically. Connector type specifies where the export should be sent. This demo environment has been preconfigured with the external ticketing system HP Service Manager. DRA can integrate with all the major external ticketing systems.

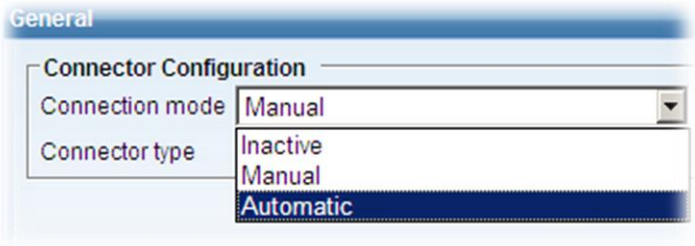

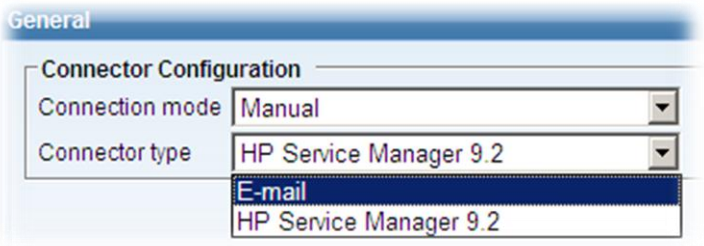

Let's configure automatic export of specific tickets to E-mail. Select **Automatic** and **E-mail** in the dialog box. Now select **Add Criteria** (❶ below).

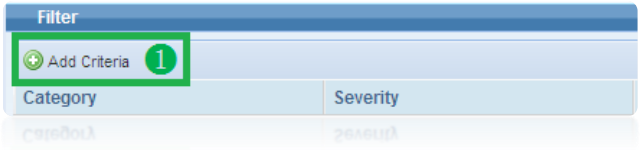

In the dialog box that appears, specify the criteria for the tickets you would like to have automatically exported.

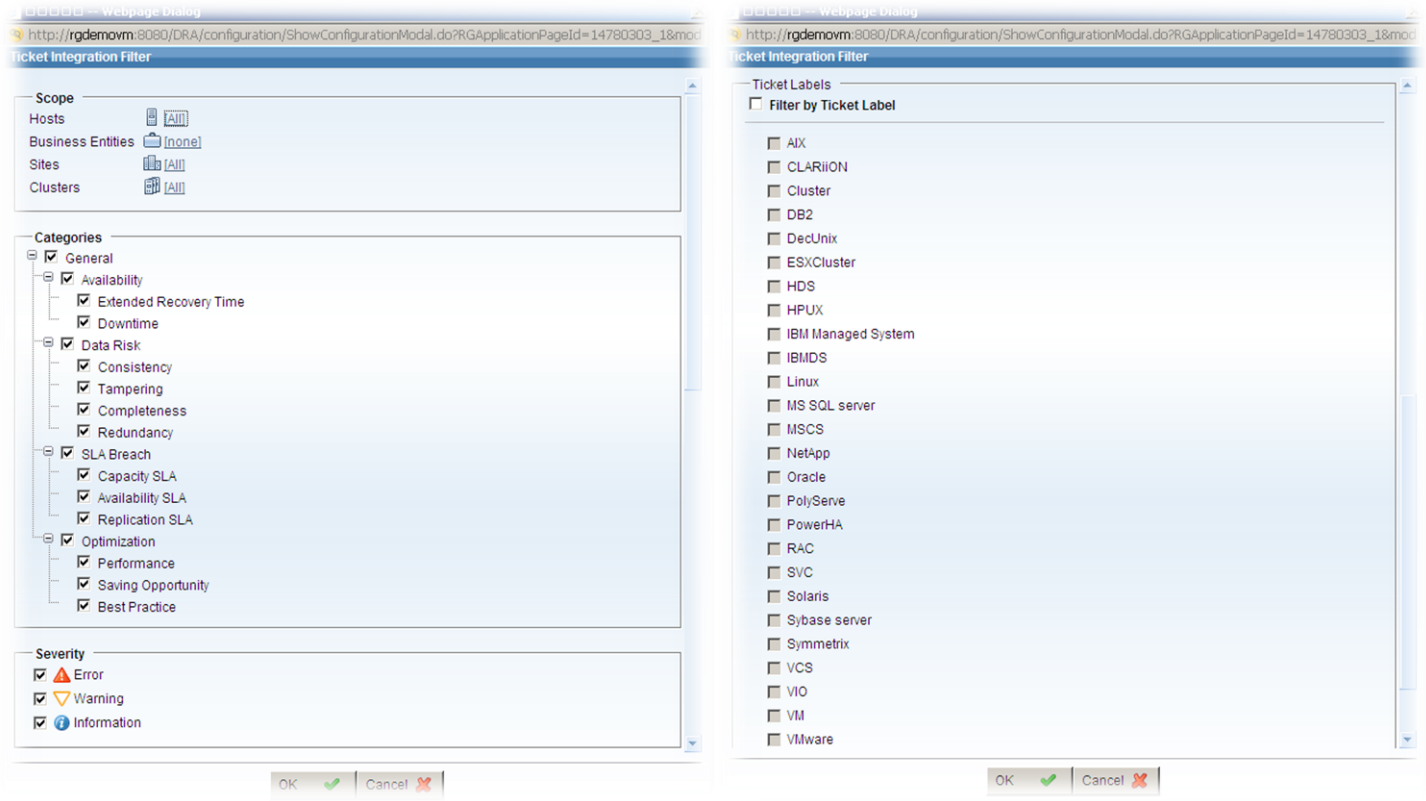

The new rule you have created now appears in the window and tickets matching your criteria are sent to the location you specified following the next scan.

Lab Exercise 7

Topic – Automatically sending newly detected risks to relevant personnel	Duration ≈ 10 minutes

The scenario for this exercise is the following:

- Each IT team must automatically receive information about newly-detected risks for their area of responsibility.
- In this exercise we will show how to schedule reports to be sent to the storage team when relevant risks are detected.

Define a Scheduled Report Task

- Select the **Configuration** tab (❶ below), and then click **Report Scheduling** (❷) on the left navigation tree.
- At the bottom of the screen, a set of buttons is available which enables users to Add, Edit, and Delete scheduled report Tasks
- Click **Add Report Task** (❸).

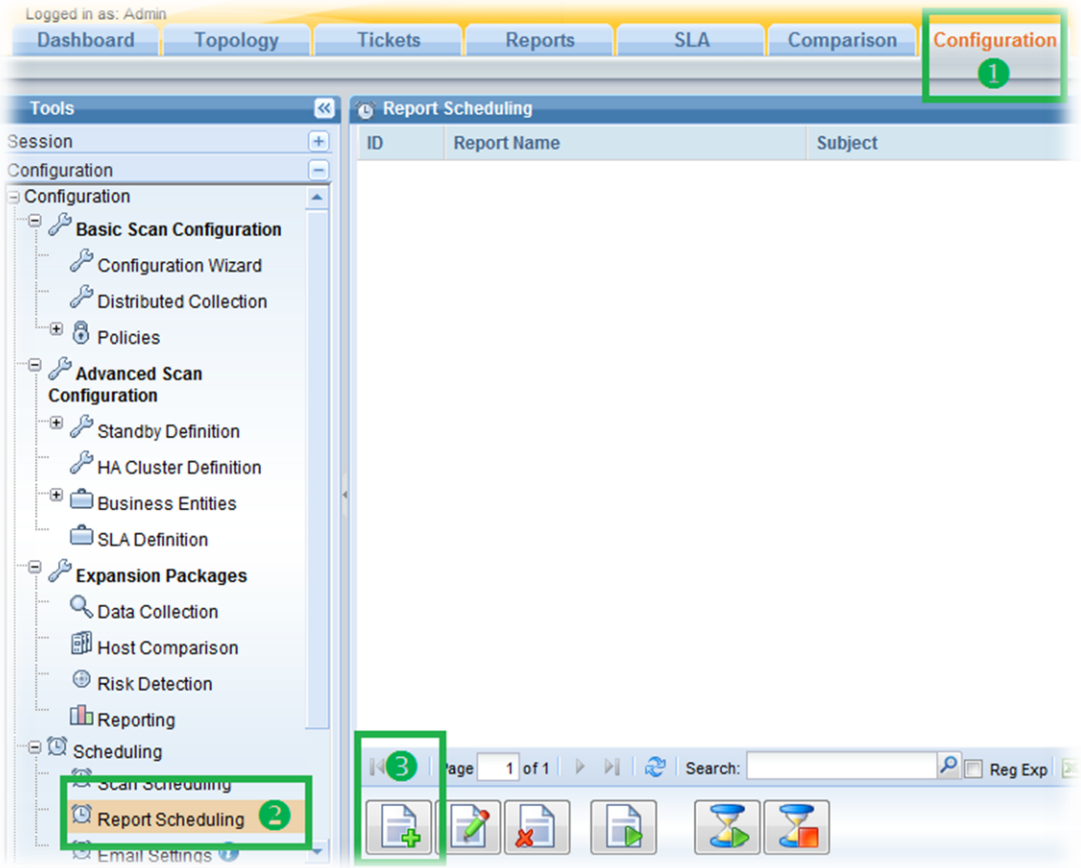

A new Window opens, presenting various reporting and scheduling options.

As a first step, select the **Ticket Details Report** (❶ below). To limit the scope of the report to newly detected issues from the last day, select **Recently opened** and **day** for the **Use Filter** (❷) and **Opened in the last** (❸) drop-down menus correspondingly.

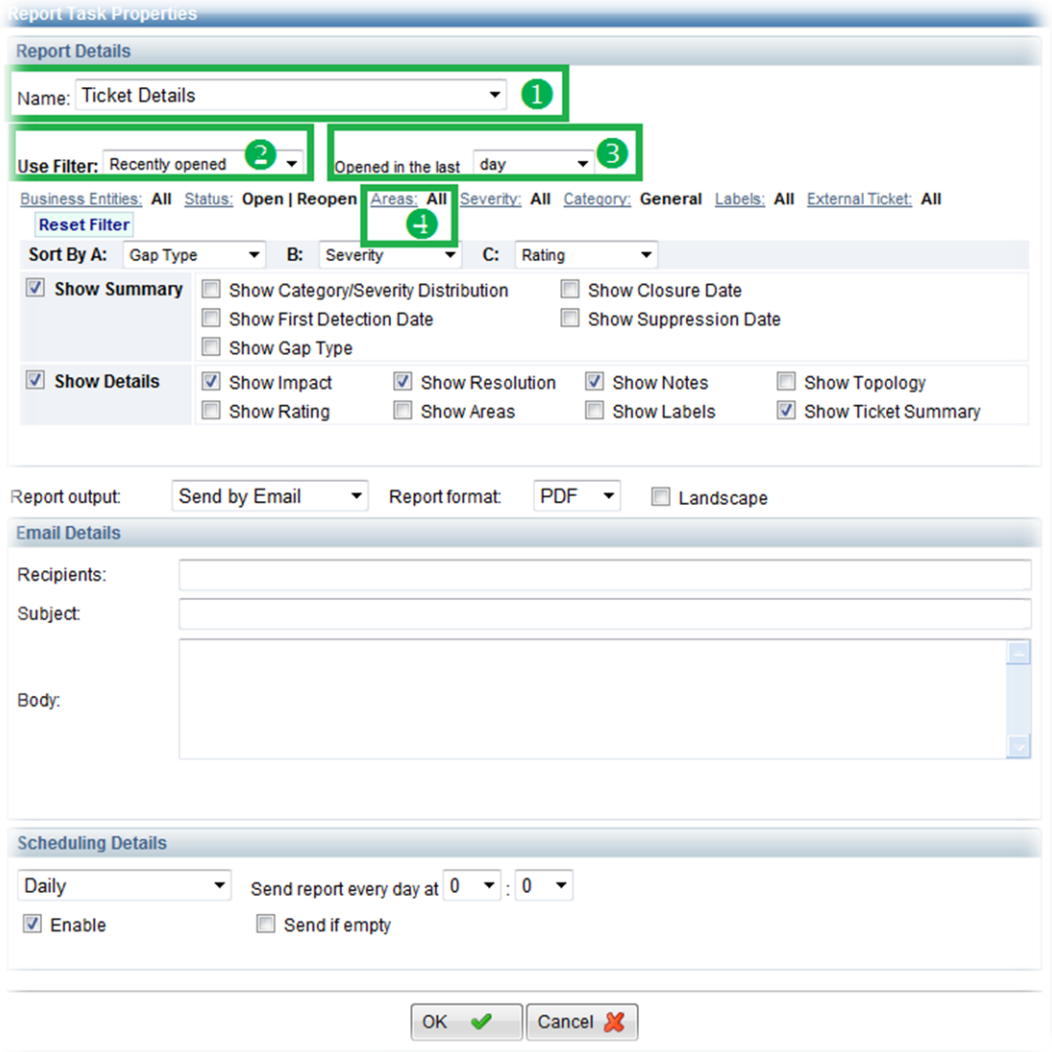

Next, to scope the report to storage-related risks only, click on the **Areas** link (4). A new window opens, presenting the different types of Areas and Labels available for ticket filtration. By default no filtration is defined. As shown in the image below, click on the **Filter by Responsibility Area** checkbox, and then click on the **Storage** checkbox below it. While not relevant for this exercise, note that you may further limit the scope to specific array types by selecting the appropriate labels (HDS, IBM DS, and so on). Click **OK** to save your changes.

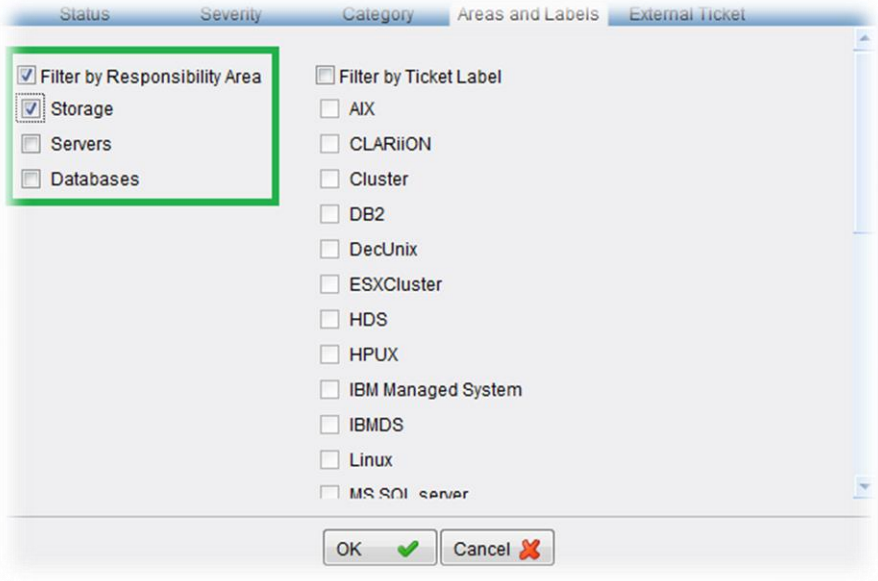

Back in the **Report Task Properties** window you can now see that the **Areas** field shows **Storage** (instead of **All**). You can define additional filtration criteria based on **business entities**, **severity**, **category** and other criteria as needed. Furthermore, the user can control what components and fields are included in both in the **Ticket Summary** and the **Ticket Details** sections of the report. In addition, using the **Report Output** option you can choose whether to send the report by email or to save it to a specific folder (local to the DRA server or a network file system), and you can also control the **Report Format** (PDF, RTF, XLS, HTML). For the purpose of this exercise, we will use the default settings; however, feel free to experiment with different options such as **Show Topology**, **Show Rating**, or any other available option.

As a final step:

- Define the **recipients** for this report (❶ below). Separate multiple recipients with commas.
- Enter the **subject** of the email (❷). You may also edit the email **body** (❸).
- Under **Scheduling Details**, select to run the report on a **daily** basis (❹) and select the required time in day (❺).
- Click **OK.**

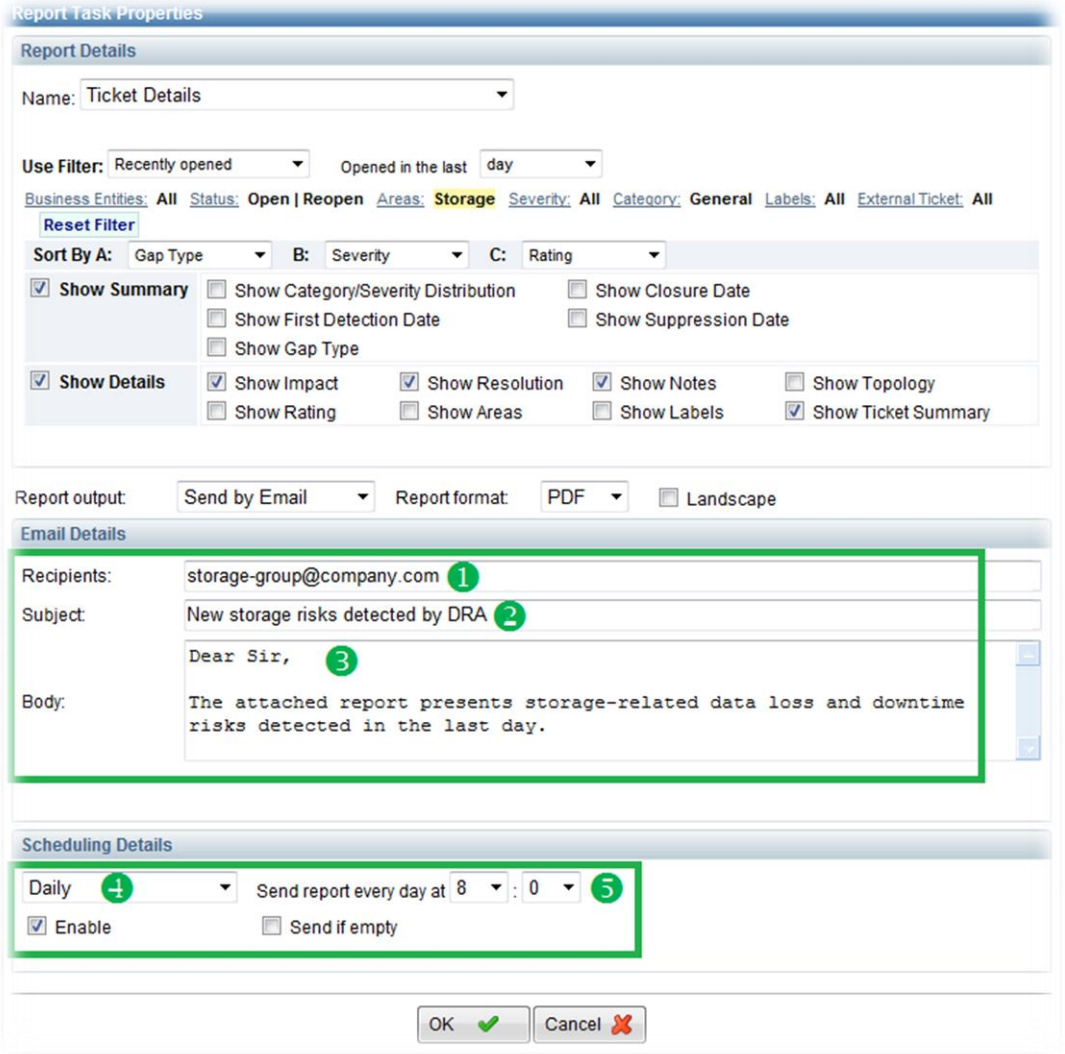

You have successfully automated routing storage risks to relevant personnel on a daily basis. Back in the main Report Scheduling page, a new Scheduled Report is now presented (❶ below). You can use **Edit Report Task** (❷) to modify the task settings or choose to run it now by clicking **Execute Now** (❸).