

# Example Application Flow Diagram

---

## Introduction

The purpose of this document is to describe a low complexity application and detail the parameters that must be obtained about that application in order to correctly configure the application for monitoring within CA Application Delivery Analysis (SuperAgent). This document also attempts to identify the types of people responsible for obtaining/providing that information about the application and infrastructure to the NetQoS administrator. Given the diversity of modern organizations, the recommended roles may not have the information required.

## XYZ Corporation's Orion Application

The fictitious XYZ Corporation's primary business application is the Orion application, which processes orders from customers. The primary interface is web based and is used by the sales people out in the field. XYZ Corp's IT department has implemented modern technology including load balancing appliances, VMWare, single-server multiple engine configurations, control port applications, and SSL.

The application infrastructure is linear as far as the data flow. There are four tiers, each dependent on the next tier. The first tier of the application is a load balancing appliance similar to an F5 BigIP. That load balancer is hosting a VIP and application port: 10.20.30.55:443, which has a FQDN of orion.xyz.com. The connection from the end users to the load balancer uses SSL, which is terminated on the load balancer. The front end network connection of the load balancer is a physical connection to the XYZ-LAX-SW1 switch on port Gi4/13. The load balancer users both destination and source IP address translation. The back end network connection of the load balancer is a physical connection to the XYZ-LAX-SW1 switch.

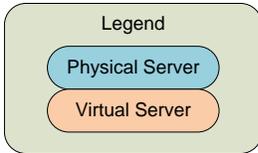
The next tier consists of two web servers in a server farm, hosted on a single VMWare ESX server. Their network connection is through port group 20 on the ESX host. The connections through the load balancer to these servers are clear http on port 80 and are identical regardless of which web server services the request.

The third tier consists of a single application server hosting three instances of the application processing engine. Each engine is bound to a unique TCP port. This application server is hosted on a different VMWare ESX server from the web server farm. The requests to the individual processing engines are identical regardless of which engine services the request. The load balancing strategy is built into a proprietary protocol between the processing engines and the web server, similar to a round-robin strategy.

The next tier of the application is a database hosted on a physical server connected to the XYZ-LAX-SW1 switch on port Gi1/12. The database application is a control port application. This means that an initial connection is made to a control port to request the connection specific parameters for a data

connection. After connecting to the control port (TCP 139), the server instructs the client to connect on a specific TCP port within the range of 5000-6000. The control session is then terminated and the data connection is established from the client to the server.

The application flow diagram for the Orion application is shown in Figure 1 on page 3.



# XYZ Corporation's Orion Application

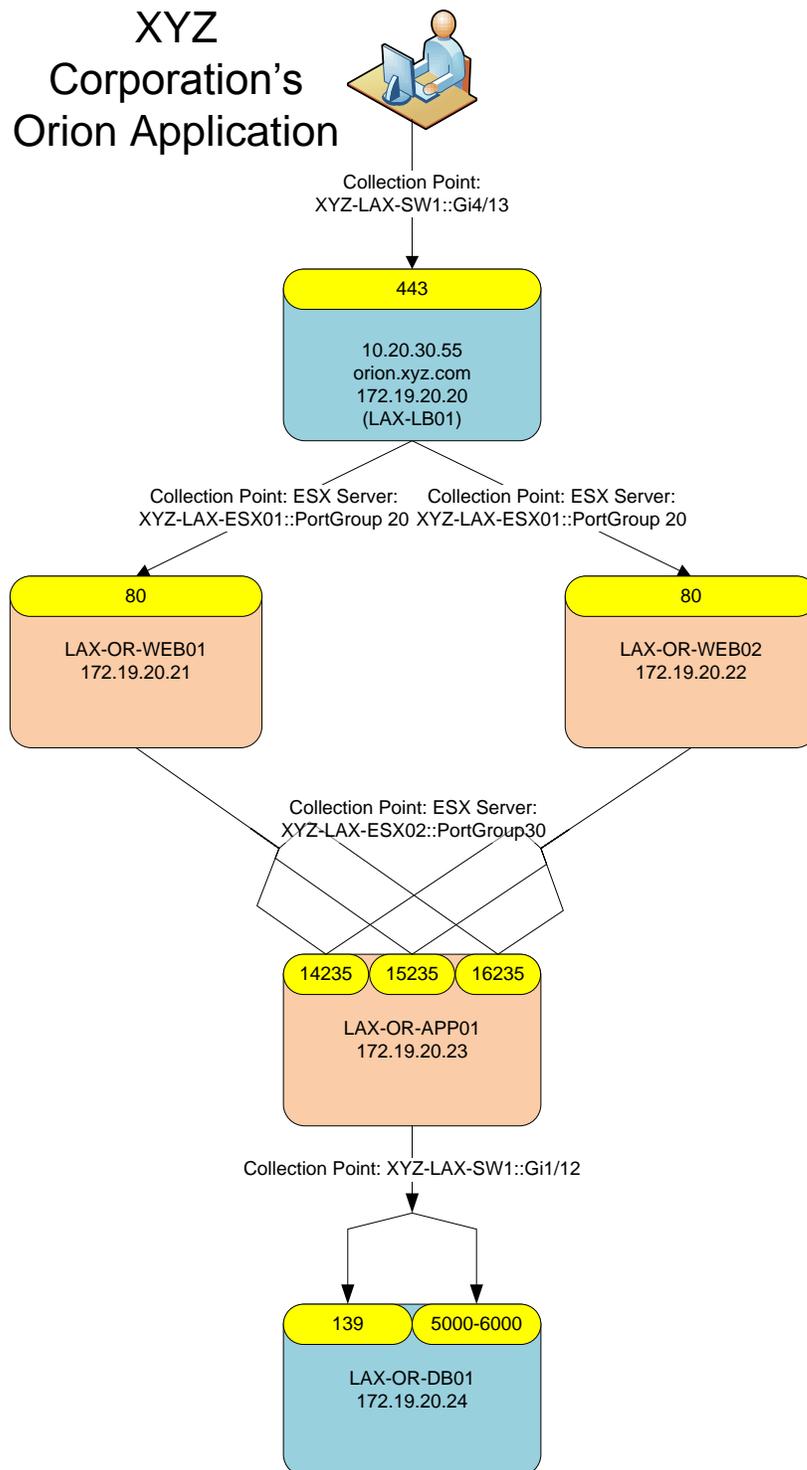


Figure 1: Orion Application Flow Diagram

## Collection Strategies

In order to monitor the Orion application through CA Application Delivery Analysis (SuperAgent), the traffic between the various tiers must be captured in the form of a live packet capture. In order to obtain the most accurate measurements, the data collection should be configured to capture the data as it goes into and out of the servers' network cards. The [CA Infrastructure Global User Community](#) hosts [details around best practices for data collection](#) in complex scenarios. These best practices should be taken into consideration when planning the data collection strategy. The collection strategy is outlined in Table 1 below.

Tier	Client	Server	Collection Point	Collector Type
1	End Users	orion.xyz.com	XYZ-LAX-SW1::Gi4/13	Physical
2	LAX-LB01	LAX-OR-WEB01 LAX-OR-WEB02	XYZ-LAX-ESX01::Port Group 20	Virtual
3	LAX-OR-WEB01 LAX-OR-WEB02	LAX-OR-APP01	XYZ-LAX-ESX02::Port Group 30	Virtual <sup>1</sup>
4	LAX-OR-APP01	LAX-OR-DB01	XYZ-LAX-SW1::Gi1/12	Physical <sup>2</sup>

Table 1: Collection Points per Tier

In the case of the physical collection, data collection aggregators may be used, as long as the aggregators do not add significant jitter to the packet arrival times at the collector.

The overall collection strategy is shown in Figure 2 on page 6.

In the case of the Orion application, the data collection for the first tier of the application should be performed at the Gi4/13 interface on the XYZ-LAX-SW1 switch instead of the core switch or a DMZ switch and should be sent to a physical collector. This could be done by way of a SPAN aggregator so long as the aggregator doesn't induce significant jitter in the captured packets.

The second tier of the Orion application could be captured at two possible points: (1) the switch port where the back end of the load balancer connects to the network or (2) the port group where the virtual web servers connect to the ESX networking stack. Of these two options, option two (2) is preferable since it provides a better vantage point for data collection<sup>†</sup>.

For the third tier, collection should be performed on a vCollector assigned to XYZ-LAX-ESX02::PortGroup30. This is a better vantage point than the vCollector assigned to XYZ-LAX-ESX01::PortGroup20 since it is closer to the server. Technically, the communication between the web servers and the app server will have to be routed via a layer 3 device since they belong to different port groups (VLANs). This means that this communication will probably reach the physical environment where it could possibly be collected without the need of a vCollector.

However, if vMotion is configured to reassign servers automatically and one vCollector is already in use to capture the tier 2 traffic, vCollectors must be deployed to every ESX host. In this case, there will

---

<sup>1</sup> This virtual collector must be distinct from the collector on XYZ-LAX-ESX01.

<sup>2</sup> This collector may be the same or a distinct collector from the collector monitoring tier 1.

already be a vCollector on XYZ-LAX-ESX02. In that case, both vCollectors on both ESX hosts will see the data and send the result up to the CA Application Delivery Analysis (SuperAgent) master console. The master console will compare the data and figure out which has the better point of view and use that data. This is the preferred mechanism for monitoring servers that can 'move' via a mechanism like vMotion.

The last tier of the application is the physical database server and would be monitored via normal physical means (a SPAN on XYZ-LAX-SW1 including Gi1/12 as a source interface).

## Contributors

In a physical collection scenario, there are two possible contributors for this type of information: (1) the server administrators or (2) the network engineer. This usually depends on the internal division of labor between the server administrator and the network engineer. If the server administrators are involved in making decisions about the connections to the network, they may have this information. More than likely however, the network engineer will be able to provide the information pertaining to the ports on a switch into which a server is connected. In a virtual environment, usually the VMWare administrator will know which ESX server and which port group are involved with a particular server or set of servers. The contributors should keep in mind that the goal of CA Application Delivery Analysis (SuperAgent) collection is to have a vantage point as close to the server NIC as possible.<sup>†</sup>

## Audience

There are two people who need this information, depending on whether the collection point is virtual or physical. If the collection point is physical, the network engineer responsible for setting up the SPAN session on the switch will use this information to ensure the proper physical ports are included in the SPAN to the collector.

If the collection point is virtual, the VMWare administrator will need this information in order to setup the virtual collector and the monitor port group for data collection. For more details and other considerations, [see chapter 14 of the CA Application Delivery Analysis \(SuperAgent\) Administrator's Guide](#). VMWare's vMotion adds a layer of complexity to this situation. If vMotion is configured to automatically migrate servers without administrator intervention, CA Application Delivery Analysis (SuperAgent) monitoring must be enabled on every ESX server involved. This means that each ESX host must have the monitor port group configured, each ESX host must have a licensed vCollector, and vCollectors should be pinned to a particular host.

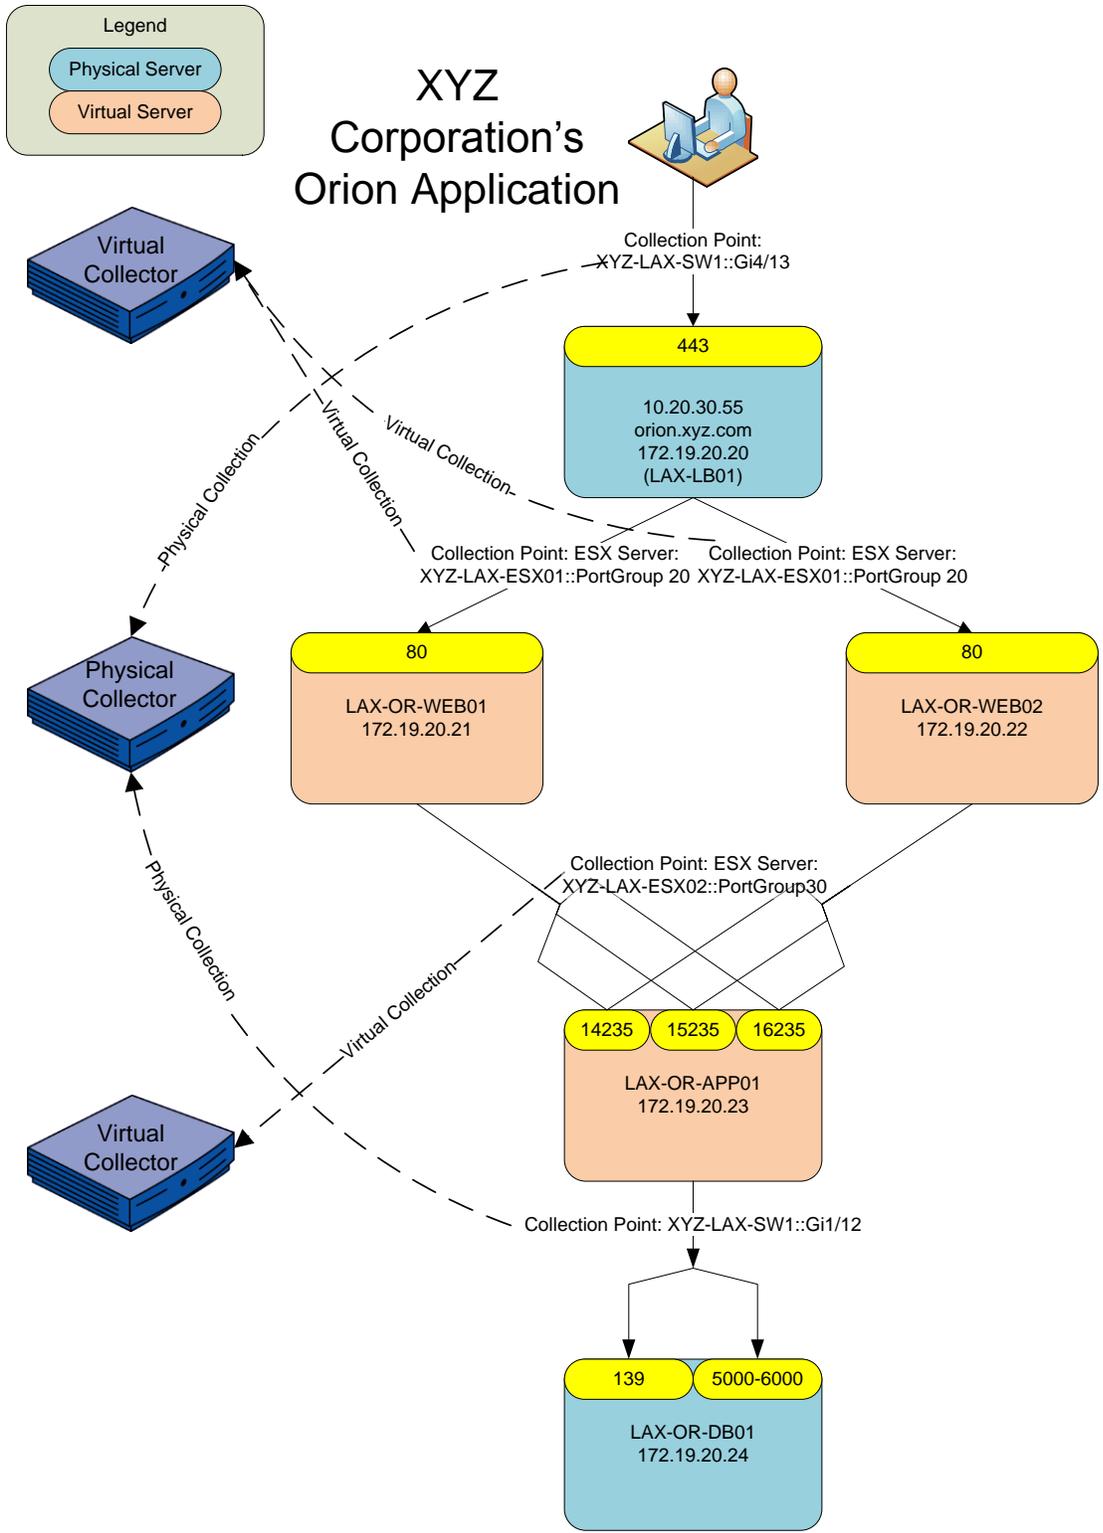


Figure 2: Orion Collection Strategies

## Application Designations

Given the infrastructure of the Orion application, five applications will be configured in CA Application Delivery Analysis (SuperAgent) using a naming standard to indicate the pertinence to the Orion application, the tier number designation, and the tier description. The applications will be configured as shown in Table 2 below and in Figure 3 on page 9.

Tier	Name	Port Begin	Port End	Servers Assigned
1	Orion - Tier 1 - Secure Web	443	443	Orion.xyz.com (10.20.30.55)
2	Orion - Tier 2 - Web Farm	80	80	LAX-OR-WEB01 (172.19.20.21) LAX-OR-WEB02 (172.19.20.22)
3	Orion - Tier 3 - Orion Processing 14235	14235	14235	LAX-OR-APP01 (172.19.20.23)
3	Orion - Tier 3 - Orion Processing 15235	15235	15235	LAX-OR-APP01 (172.19.20.23)
3	Orion - Tier 3 - Orion Processing 16235	16235	16235	LAX-OR-APP01 (172.19.20.23)
4	Orion - Tier 4 - DB Control	139	139	LAX-OR-DB01 (172.19.20.24)
4	Orion - Tier 4 - DB Data	5000	6000	LAX-OR-DB01 (172.19.20.24)

Table 2: Orion Application Designations

The first two tiers of the application are simply designated by their corresponding name and port. Tier 2 can be configured as a server farm listing a particular number of servers that should be available in order for CA Application Delivery Analysis (SuperAgent) to consider the application available. In some situations, the loss of 50% of the servers in a server farm will not crash the application. In other scenarios, losing a single server in the server farm will constitute an application outage.

The third tier of the application involves three engines hosted on the same server. Since they use distinct non-contiguous ports, the CA Application Delivery Analysis (SuperAgent) administrator would configure each engine as its own application in CA Application Delivery Analysis (SuperAgent). In most scenarios this will be more desirable.

The other option would be to make a concession in monitoring this tier of the application. If the three engines are to be monitored as one set of load balanced engines, the CA Application Delivery Analysis (SuperAgent) administrator must specify a single range of ports including the three hosted by the engines. This means that any traffic hosted by that server on ports in that range would be included in this application bucket. The obvious difference would be that the three engines would report as an aggregate server farm instead of individually.

The last tier actually involves two different services and is configured as such. Since the control service can point to any port within the range of 5000-6000, the DB Data application should be configured as a range.

## Contributors

The application owners and sometimes the vendor will be able to provide the information about the application designations. In particular each usually has a contact or point person who can provide this information.

For the Orion application, the first tier is between the users and the load balancer. The load balancer administrator (which is sometimes the same as the network engineer) will have the information about which VIPs are hosted and which virtual servers are hosted on each VIP.

This administrator will also have knowledge concerning whether or not the load balancer uses source and destination IP address translation or just destination address translation. Either scenario is valid, however when a load balancer uses only destination address translation, the packet captured between the load balancer and the load balanced servers will maintain the original user's IP address. For most monitoring tools, this means that the conversations captured between the load balancer and the load balanced servers will appear as though they are actually happening between the end user and the load balanced servers. While this is perfectly acceptable, it does require a little more knowledge to interpret the resulting data.

## **Audience**

The main audience for this information is the CA Application Delivery Analysis (SuperAgent) administrator. This administrator will take the application information and input it directly into the CA Application Delivery Analysis (SuperAgent) configuration, enabling the application's various components to be properly labeled and segregated into their own reporting groups.

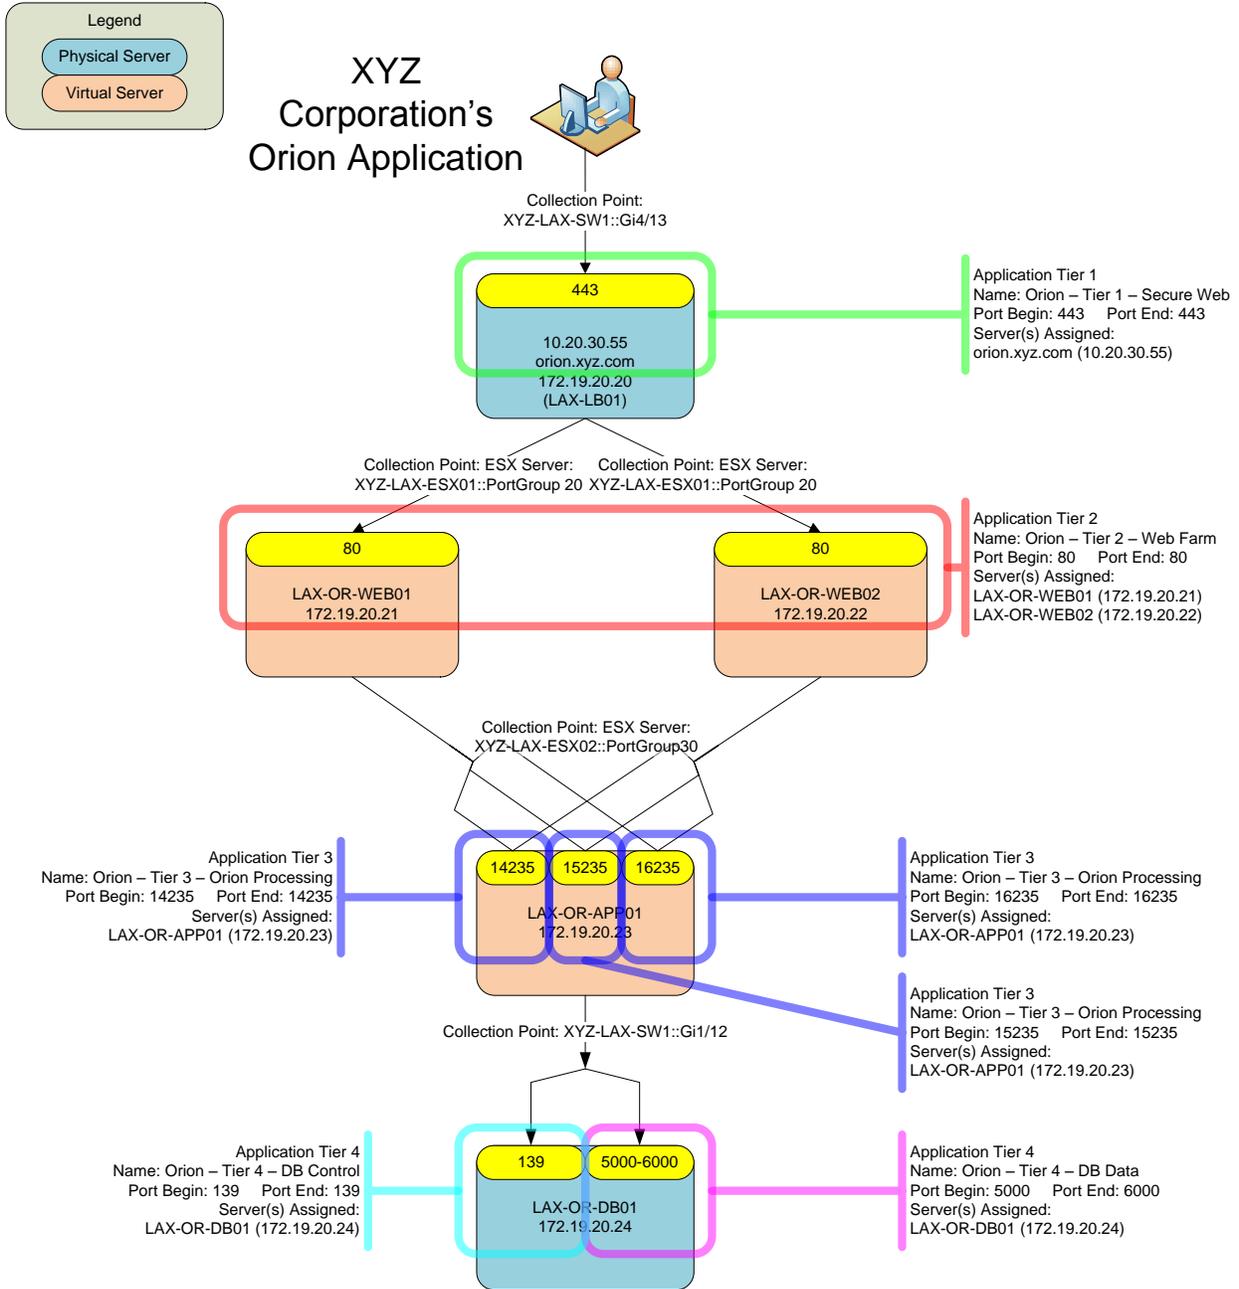


Figure 3: Application Designations

† If monitoring were performed from a less than optimal point of view, some of the metrics which usually indicate server-only induced latency would instead indicate latency induced by the server *or* the network equipment between the server and the collection point. While this is still helpful, it doesn't narrow the fault domain enough to signal root cause.