Oracle Traffic Director

Technical Deep Dive and Deployment Best Practices

Sriram Natarajan, Product Management – Web Tier, Traffic Director

ORACLE

Safe Harbor Statement

The preceding is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

Traffic Director 11g – Technical Deep Dive

- Big Picture View
- Terminology Overview, Topology Discussions
- Feature Deep Dive
- Deployment Best Practices
- Peek at Future Investments
- Q & A

Traffic Director 11g – Technical Deep Dive

- Big Picture View
- Terminology Overview, Topology Discussions
- Feature Deep Dive
- Deployment Best Practices
- Peek at Future Investments
- Q & A

Introducing Oracle Traffic Director (OTD) 11g

High performance, low overhead Application Delivery Controller

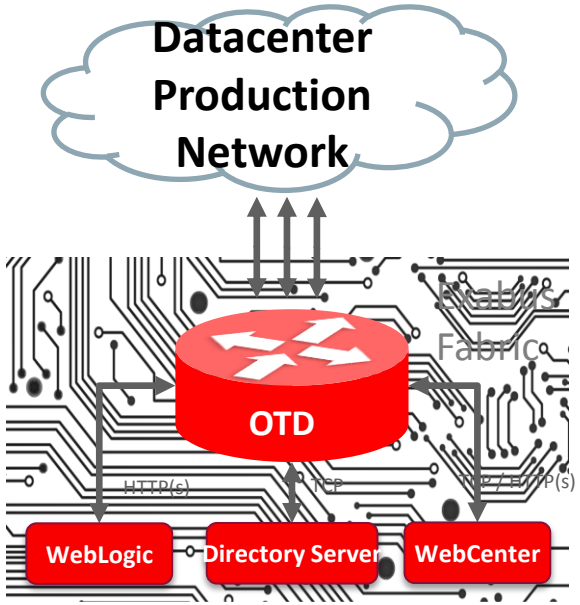

Traffic Management

High Performance

High Availability

Security

Traffic Management

- Setup HTTP/TCP Load Balancing Policies
  - Full fledged HTTP/WebSocket Reverse Proxy
    - Multiple Load Balancing Algorithms
    - Deep Health check
- Setup traffic shaping policies via
  - UI Wizards or If/Else syntax scripting
- Configure Quality of Service
  - For Incoming and Outgoing requests
    - Drain/Limit/Throttle Connection/Request

High Availability

- No Single Point of Failure
  - Node level redundancy
    - Farm deployment/Administration via Admin
    - Built-in Active/Passive HA on engineered systems
  - Watchdog for process level redundancy
- Increased Application Availability
  - Deep Health check for intelligent routing
    - Mark Application Server as Backup
  - Dynamic reconfiguration
    - Update traffic policy with no downtime!

High Performance

- High throughput; Low response time via
  - Leverage Caching/Compression in OTD
  - SSL offloading on Intel/SPARC
- OOB outperforms Apache HTTP Server in
  - Static content serving and Reverse Proxy

Security

- OOB Single Sign-On via WebGate/OAM
- Web Application Firewall (ModSecurity)
- TLS 1.2, Request Limiting (DoS Protection)
- Monitor via Enterprise Manager Cloud Control

Traffic Director 11g – Comparing with Oracle HTTP Server

Oracle Traffic Director 11g

- Application Delivery Controller (ADC)
  - Software Load Balancing with Built-In High Availability (No SPoF)
  - Integration with Engineered systems infrastructure (SDP / IPoIB)
  - Request Limiting/Throttling (QoS) with offline deep Health Check
  - Built-In WebLogic Plug-In support to front-end WebLogic servers
  - Out of the box configuration replication, distributed administration
  - Extremely light weight, wizard based administration, CLI
  - Dynamic configuration, minimal server restart
  - Certified with FMW, PeopleSoft. Supports Load Balancing with EBS
- Licensable via
  - EECS on engineered systems (Exalogic/SuperCluster)
  - Oracle Access Management (Suite Plus) , Oracle Single Sign-On (Suite Plus)
  - WebLogic Software Development Kit for Database Appliance

Oracle HTTP Server

- Primarily Web Server with Reverse Proxy capabilities
  - Certified with FMW, PeopleSoft, Siebel, EBS stack
  - Supports content serving (HTML, Images, CGI/Fast CGI)
  - Includes WebLogic Proxy Plug-In to front-end WebLogic
  - FIPS 140-2 compliance
  - Management via WebLogic management framework
  - Some configuration changes applied via Soft Restart
    - Some impact on persistent connections
- Licensable via Web Tier SKU
  - Included within WebLogic license to host on same WLS processor

Traffic Director 11g – Release Lifecycle

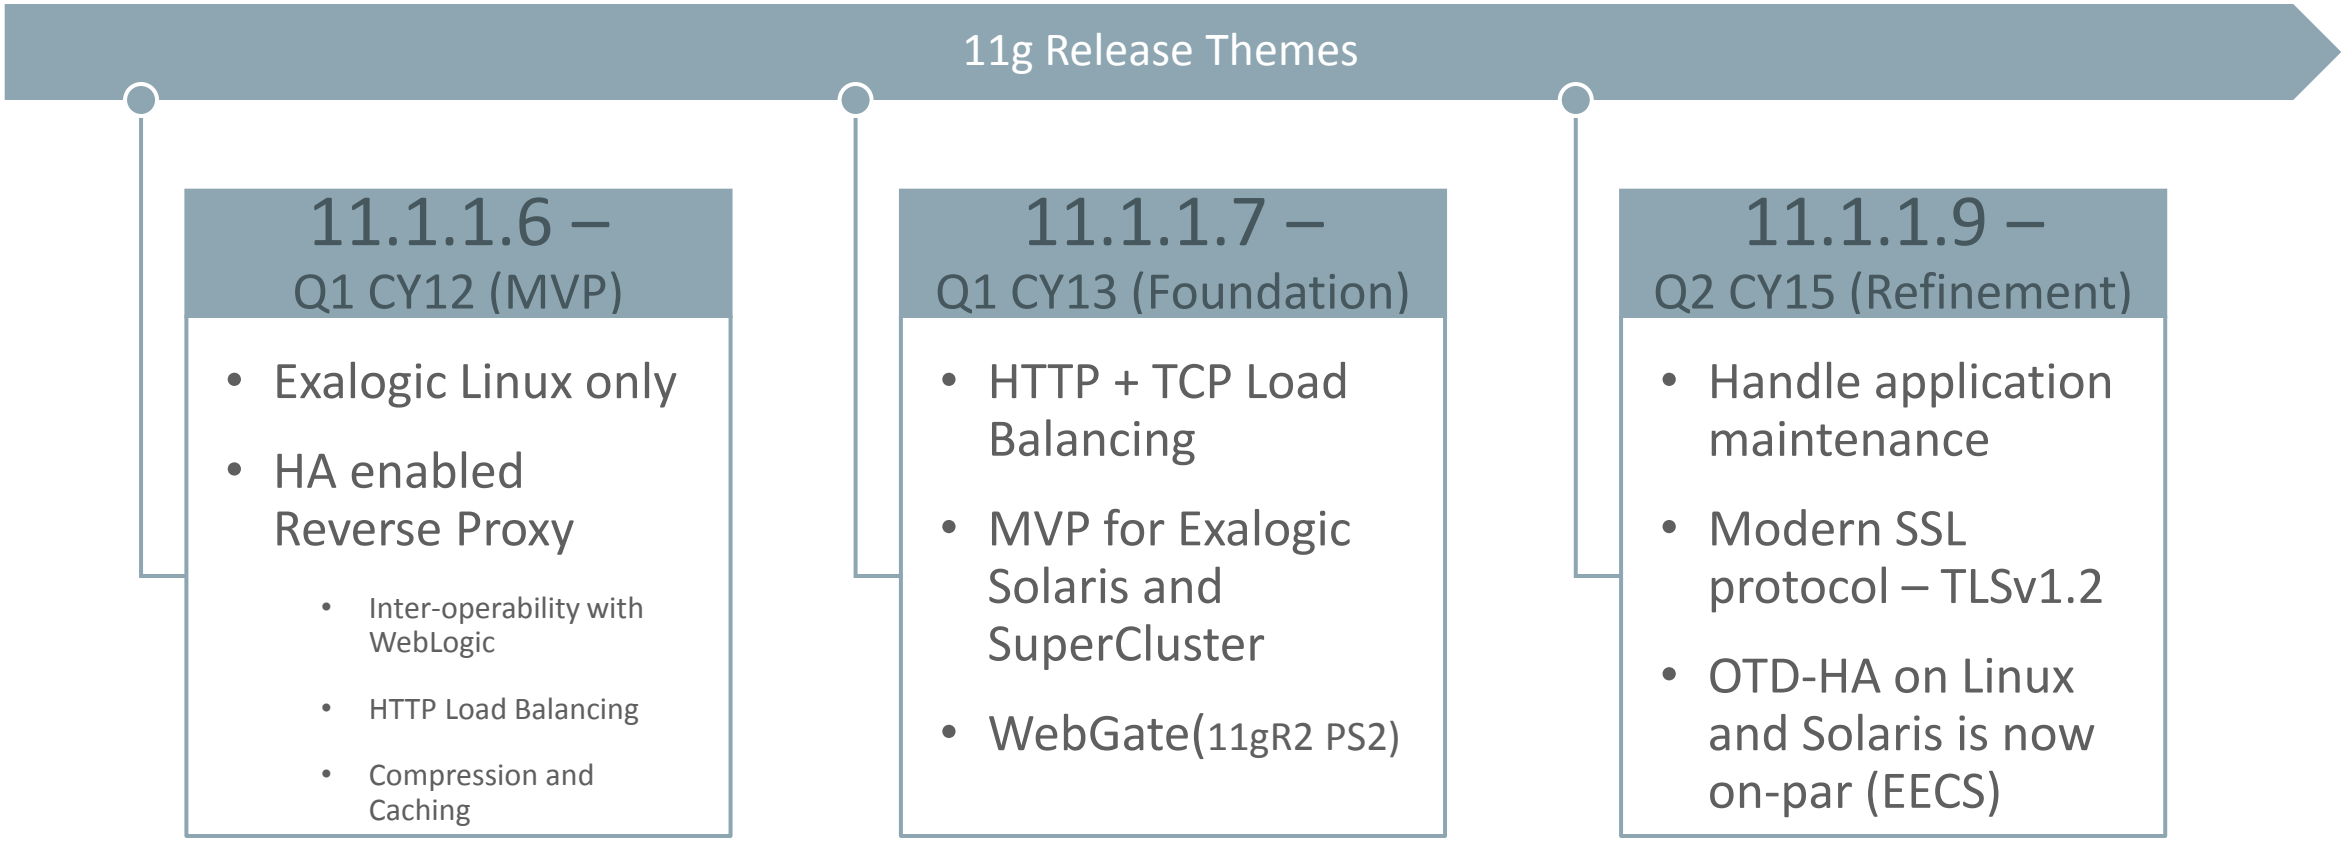

Traffic Director 11g – Platform, Patches and Support

- Supported Platforms
  - Oracle Enterprise Linux 5.6+ and Oracle Enterprise Linux 6.5+
    - Implicitly supports RedHat Enterprise Linux 6.5+ (for Oracle Access Management use cases)
  - Solaris 11.1+
    - 11.1.1.9 requires Solaris 11.2+ on engineered systems
- Patches
  - Integrates with Oracle Patch infrastructure; Standard Oracle patch policies apply
- Support Dates

Product	GA Date	Premier Support Ends	Extended Support Ends	Sustaining Support Ends
Traffic Director 11g	Mar 2012	Dec 2018	Dec 2021	Indefinite

Traffic Director 11g – Technical Deep Dive

- Big Picture View
- Terminology Overview, Topology Discussions
- Feature Deep Dive
- Deployment Best Practices
- Peek at Future Investments
- Q & A

Traffic Director 11g – Key Concepts

Administration Server

- Close to WLS Admin Server
  - UI (browser based) and CLI
- Manage config and server lifecycle on multiple machines
- Acts as 'Admin Node' on same machine!

Administration Node(s)

- Close to WLS Node Manager
  - Registers with Admin Server via SSL
  - Executes Admin commands on local machine (incl. server lifecycle Mgmt.)

Configuration

- Blueprint (Abstract) Includes:
  - External End Point (Listener)
  - Traffic Handler (Virtual Server/TCP Proxy)
  - Traffic Shaper (Routing Rules)
  - Origin Server (back-end)
  - Deployed on Admin Server and/or Node(s)

Load Balancer Service

- Deploy configuration to Admin Node
  - Operates independent of Admin
  - Actual config files, processes
  - Separate from Admin
  - Includes server lifecycle Mgmt.
  - Also known as Server Instance

Failover Group

- Built-In High Availability
  - Deploy Configuration on 2 Admin Node(s) – Becomes LB services !
  - VIP (Floating IP) front-ends this Load Balancer service running on 2 Admin Node(s)
  - No SPoF even when a machine crashes!
  - Single VIP – Active – Passive HA
    - LB service 'Primary' on only 1 Admin Node
    - Backup takes over only when Primary does not respond
  - VIP Pair - Active – Active HA
    - Requires external DNS load balancing
  - Limited to engineered systems

Traffic Director 11g – Development oriented topology

Simple to provision, administer

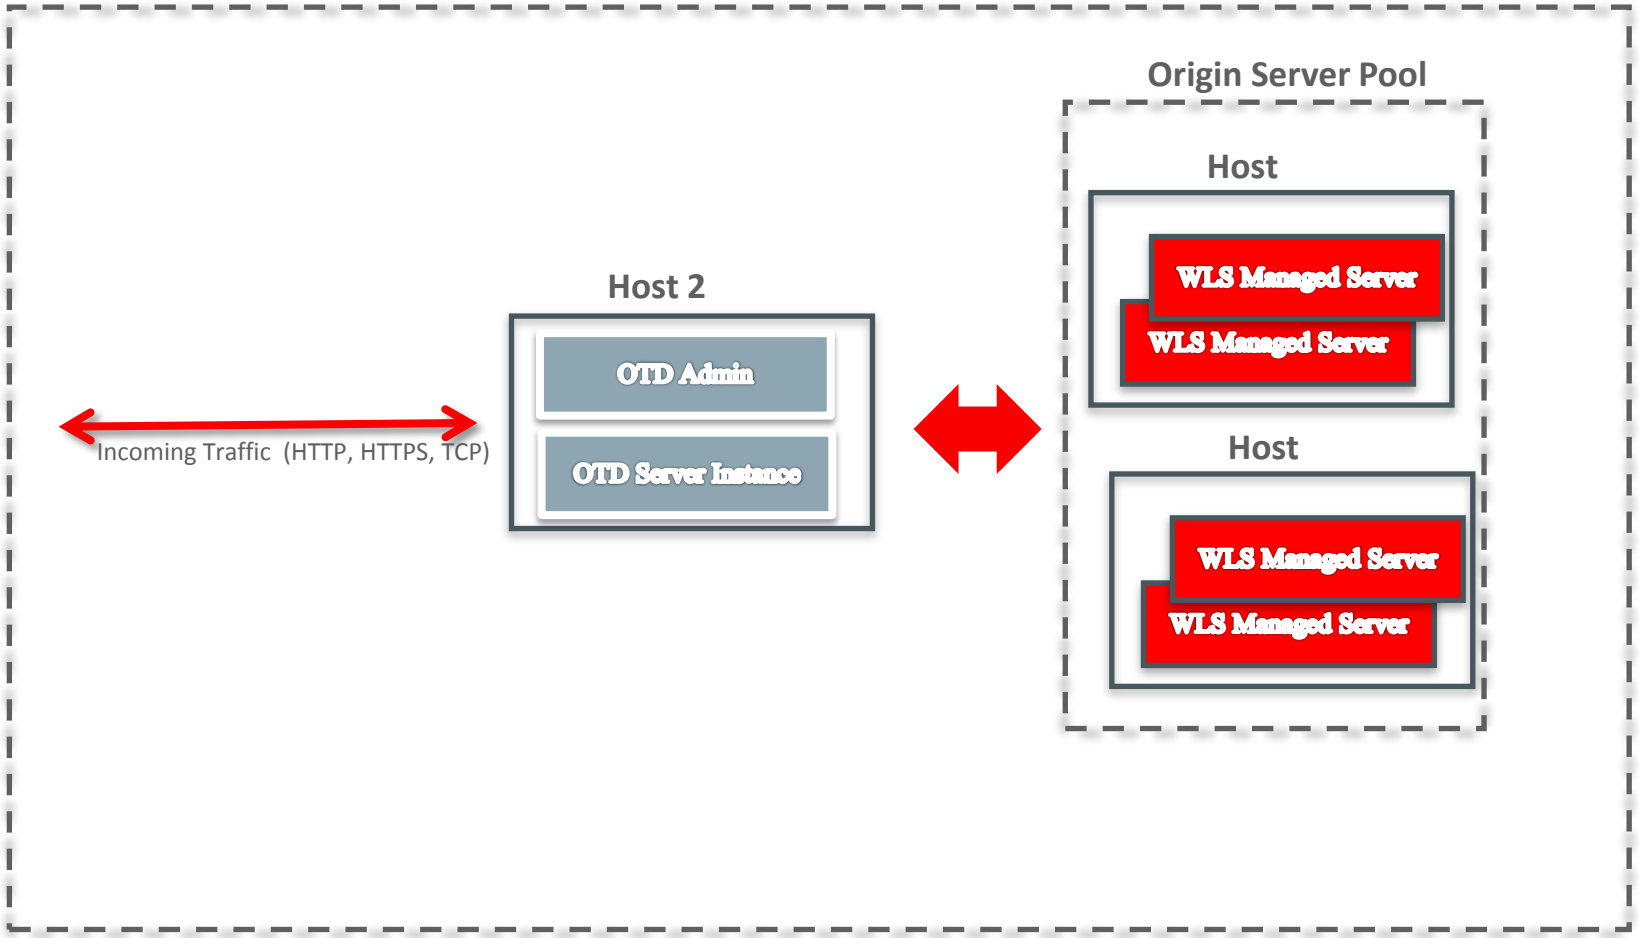

- Install OTD
- Configure OTD Admin Server
  - Default listens at port 8989
- Create OTD Configuration
  - Choose HTTP/TCP Load Balancing
  - Default listens at port 8080
  - 'Deploy Config'
    - Translate to OTD Server Instance on all machines
    - Does the actual Load Balancing
- Automate via CLI

Traffic Director 11g – Production oriented topology

Increases overall application availability

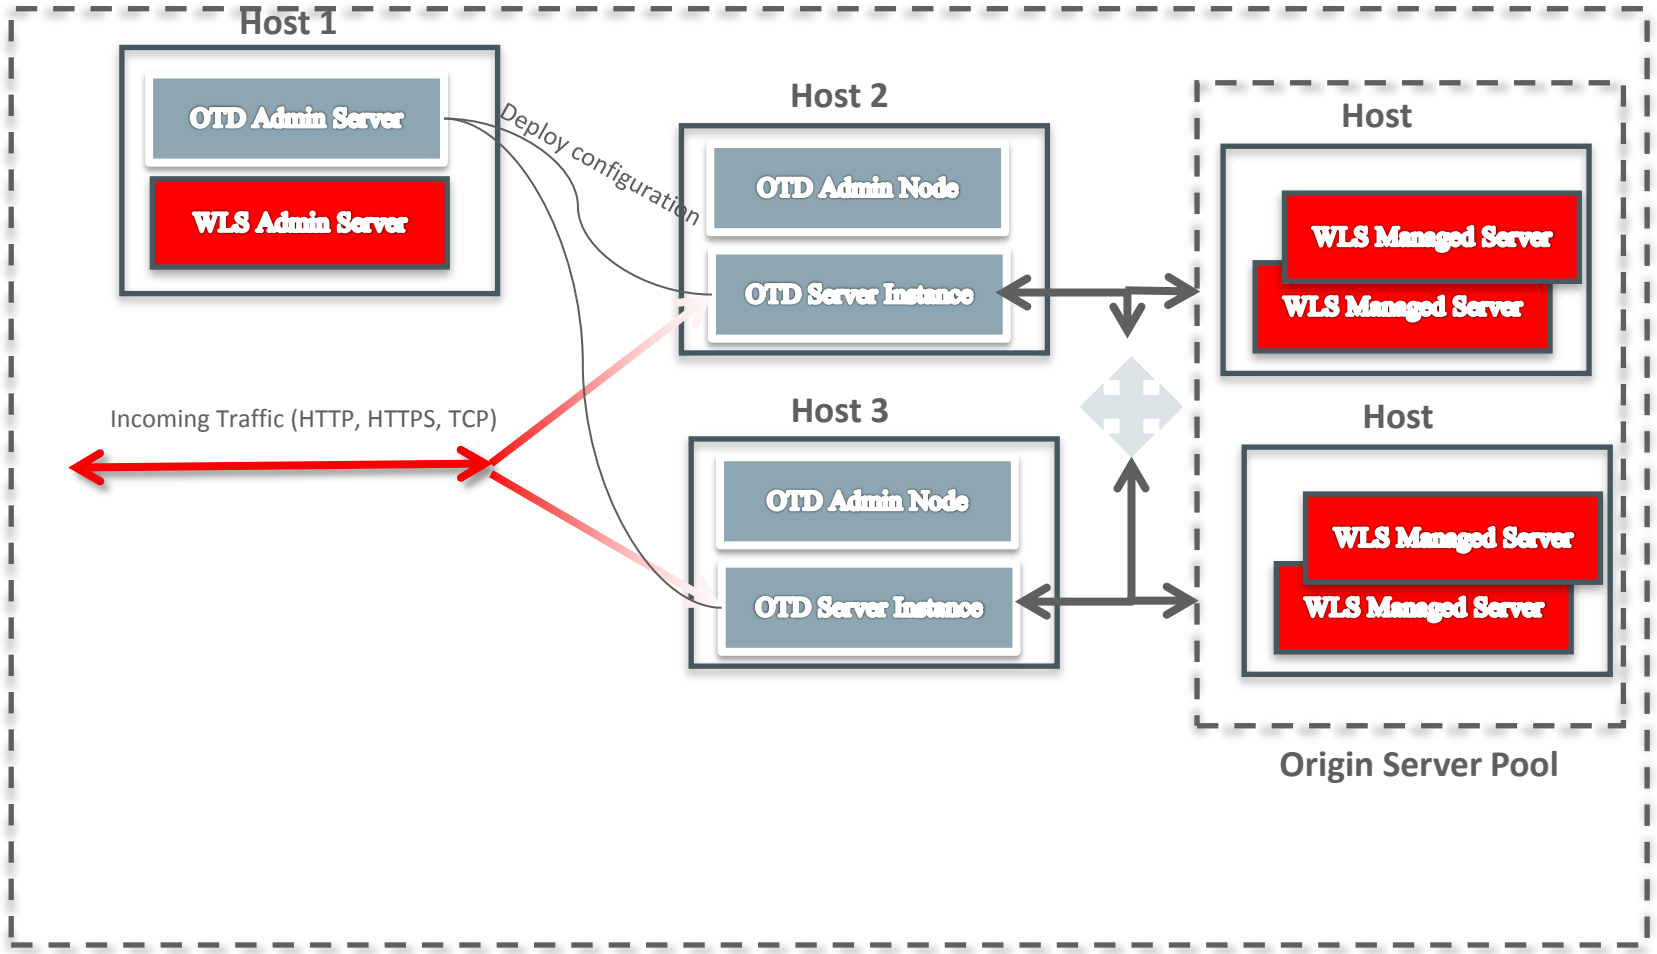

- Install OTD on Host 1,2,3
- Host1: OTD Admin Server
  - Single place to administer OTD farm
  - Default listens at port 8989
  - Create OTD Configuration
    - Choose HTTP/TCP Load Balancing
  - Translate to OTD Server Instance
    - Does actual Load Balancing
  - Default listens at port 8080
  - Deploy configuration publishes configuration to Host 2, 3
- Host 2, 3 : OTD Admin Node
  - Pairs with **remote** Admin Server
- Complete CLI automation

Traffic Director 11g – Technical Deep Dive

- Big Picture View
- Terminology Overview, Topology Discussions
- Feature Deep Dive
- Deployment Best Practices
- Peek at Future Investments
- Q & A

Traffic Director 11g – Configuring Administration Server

- Assumes Oracle Traffic Director 11g binaries are successfully insalled
- Acts as 'Administration Node' on the local machine

Create and start an admin server:

```
cd <OTD_HOME>/bin
./tadm configure-server --user=admin
   --instance-home=/u01/app/otd/admin1

/u01/app/otd/admin1/admin-server/bin/startserv
```

By default, admin server/nodes bind to all available network interfaces and on port 8989.

Prompts for a password

Admin Server acts as Admin Node on local machine

Traffic Director 11g – Configuring Administration Node

- Assumes
  - Traffic Director 11g binaries are successfully installed
  - Traffic Director 11g administration server is configured and running
- Pairs with remote Administration 'Server' and executes on local machine

Create an admin node, register it with the remote admin server, and start it:

```
cd <OTD_HOME>/bin
./tadm configure-server --admin-node --host=myadminhost
   --user=admin --instance-home=/u01/app/otd/node1

/u01/app/otd/node1/admin-server/bin/startserv
```

Location of remote admin server

Traffic Director 11g – Administration Architecture

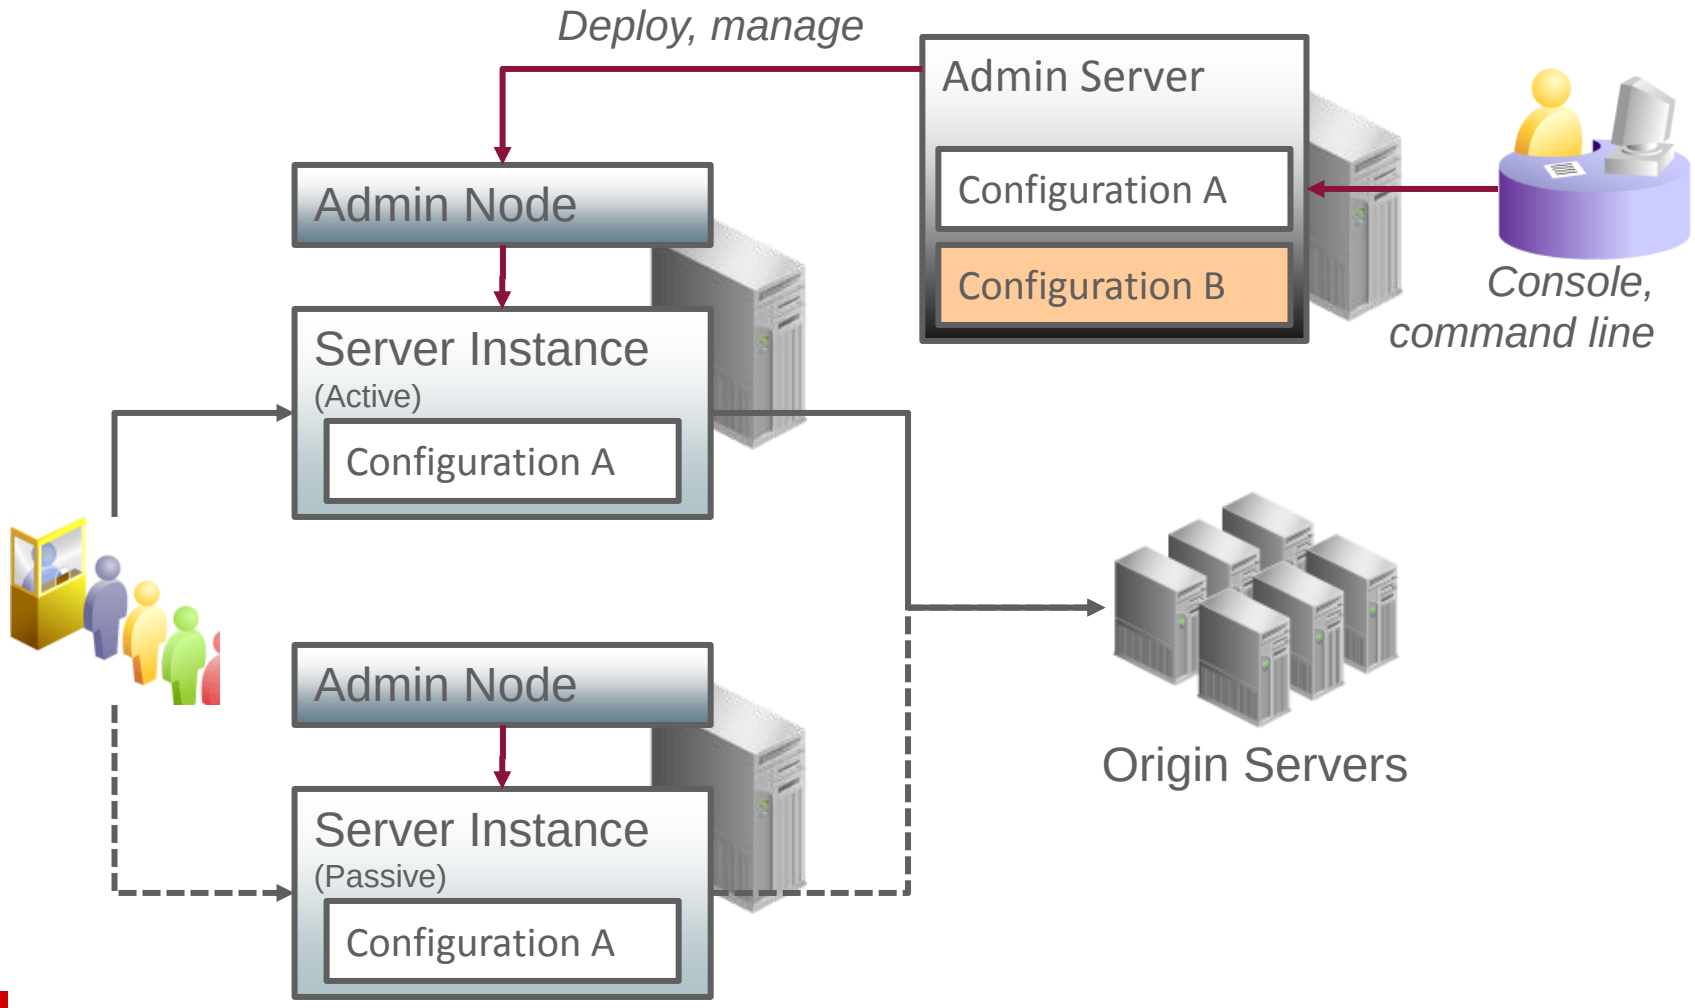

Traffic Director 11g – Admin Overview

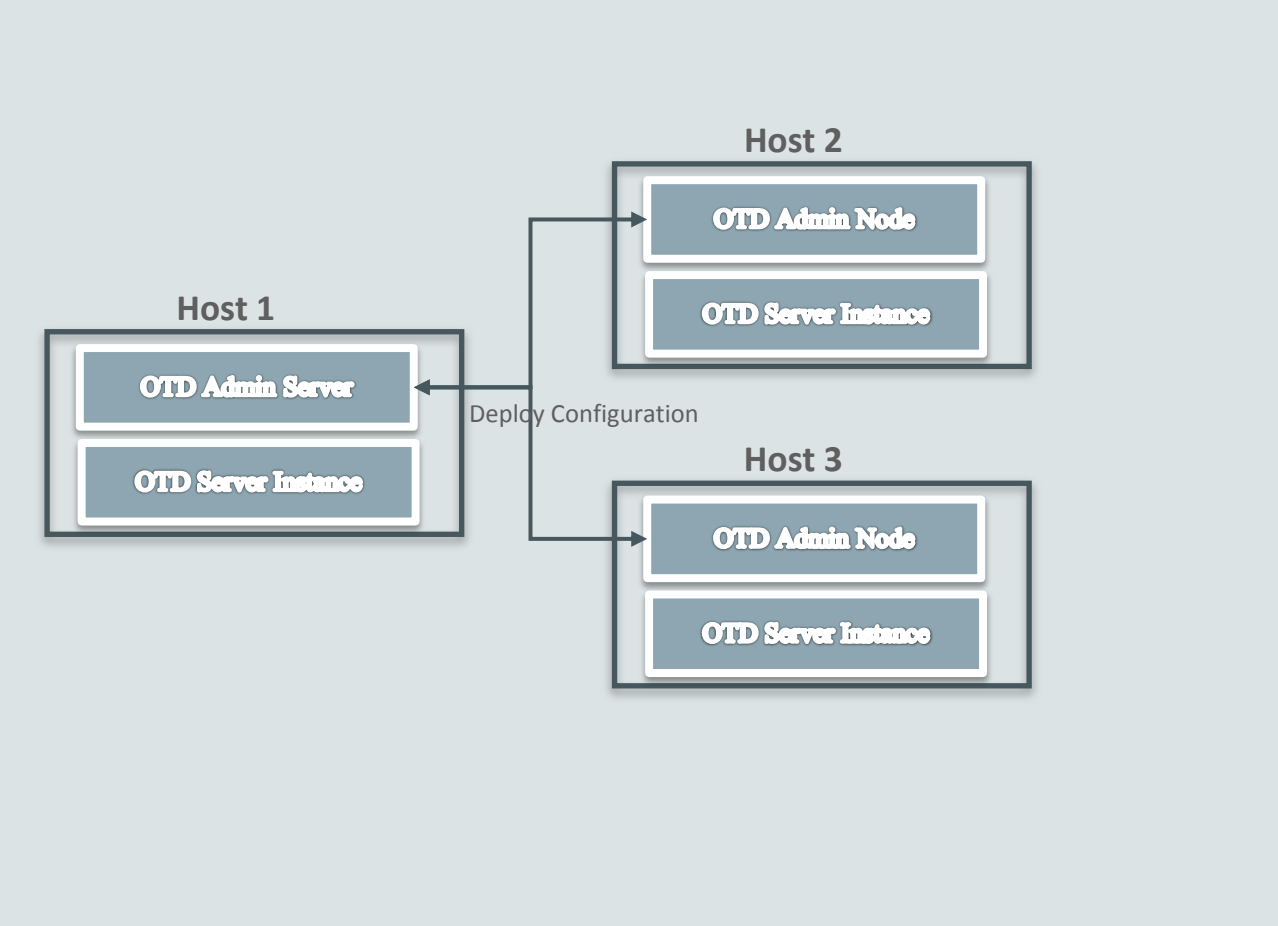

- Administration (as distributed cluster)
  - *1 Admin server* manages a OTD farm (cluster) via *Admin Nodes*
  - *Admin Server* communicates with *Admin Node(s)* via HTTPS
  - Homogenous environment
- Configuration Administration
  - 1 Admin Server manages *Configuration* on all Admin Nodes
  - Admin Server publishes configuration changes to all machines
  - Automatic Backup of configuration changes.
    - Restoring to one of the last 6 changes – extremely simple
- Server Lifecycle administration
  - *Server Instance* (based on a config or LB policy) on a machine delivers Load Balancing functionality
  - Server Instance is separate from Admin; Does not require Admin running
- Monitoring
  - Server Instance provides monitoring statistics via text/XML/SNMP
  - Admin Server collects these statistics via Admin Node(s)
    - EM Cloud Control consumes these monitoring statistics (via OTD Admin Server)

Traffic Director 11g – Configuration Architecture

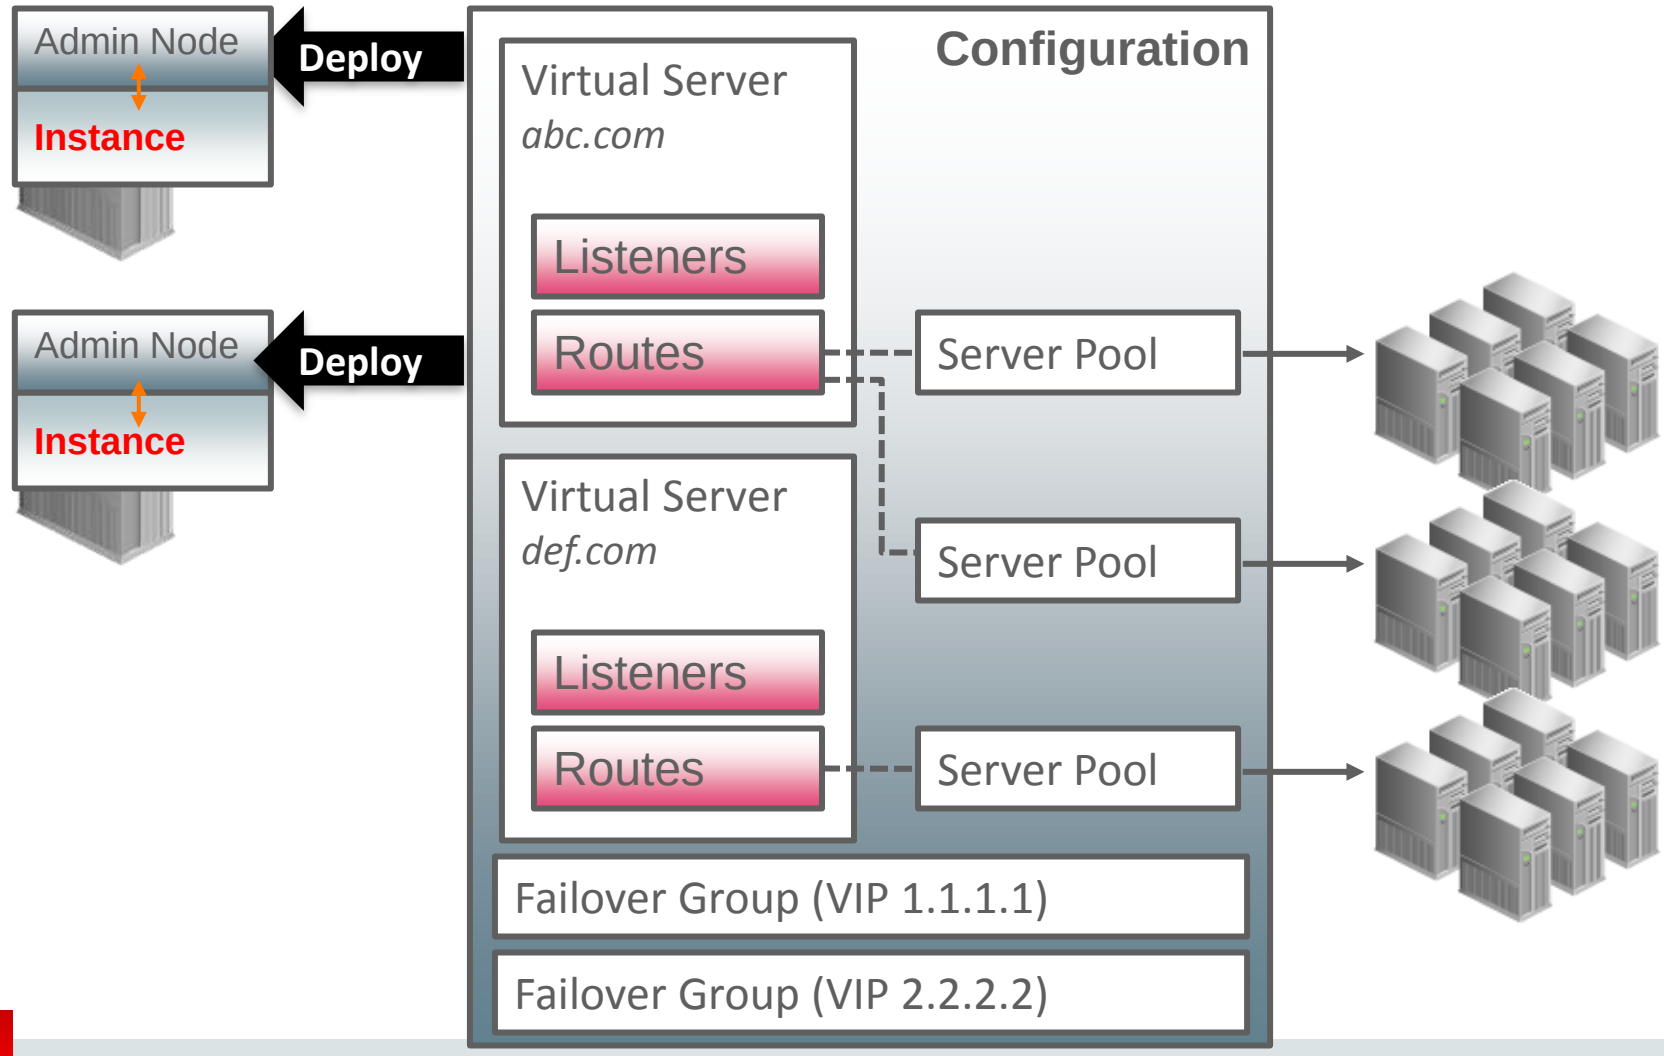

Traffic Director 11g – HTTP Load Balancing Overview

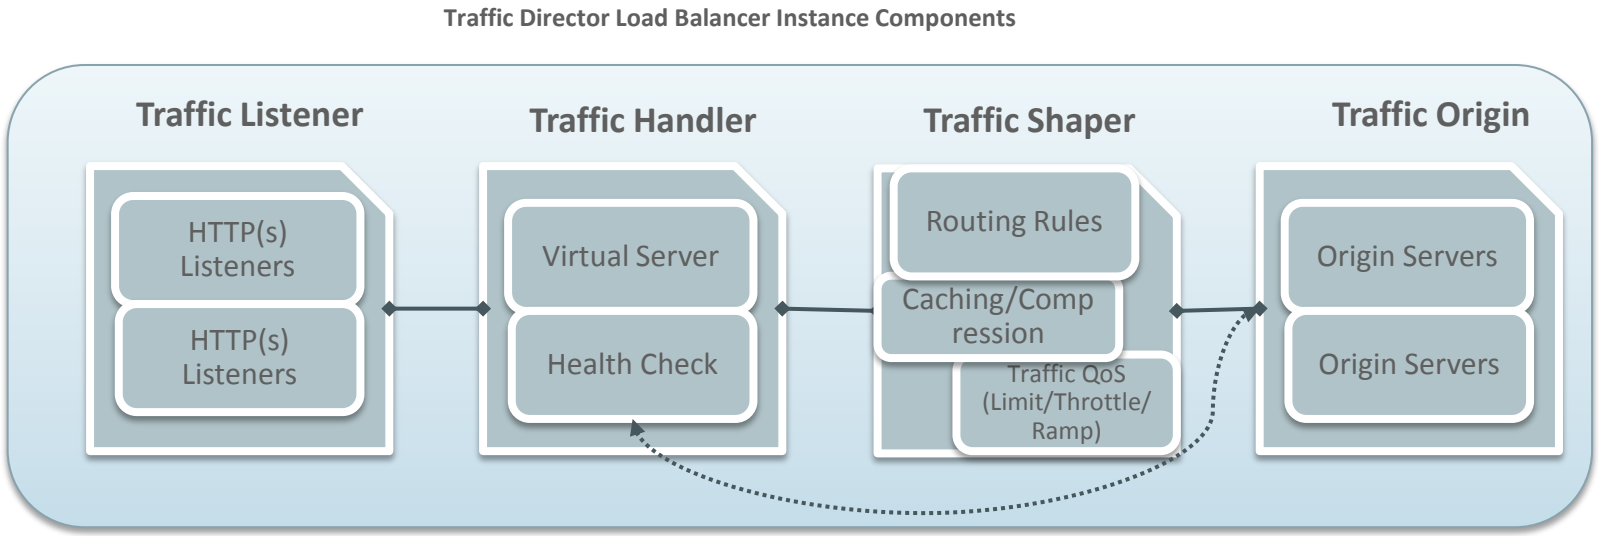

- Traffic Director HTTP Load Balancer Server has 4 key components:
  - Traffic Listener - Handles incoming connections (IP/VIP : Port)
  - Traffic Handler - Processes incoming client TCP/IP connections (no data processing)
  - Traffic Shaper - Shapes incoming traffic and can offer QoS for underlying traffic
  - Traffic Origin - Passes traffic to back-end application server where content originates

- Acceptor and Keep Alive threads handle Traffic Listener load
  - SSL Termination
    - Offload cipher to processors
- HTTP Thread Pools handle Traffic Handler/Shaper processing load
  - Traffic Routing based on any HTTP request header/body
  - Origin Server Pool selected based on Traffic Routing rules
  - Caching/Compression support on origin server response
  - QoS (Connection Limit/Ramp Up) on traffic to origin servers
  - Request limiting to protect from DoS attack
  - Serve custom HTML page when origin server is offline
- Customizable Health Check
  - HTTP GET/OPTIONS on a URI
  - Validate on response header/body
    - Deep Health Check support

Traffic Director 11g – TCP Load Balancing Overview

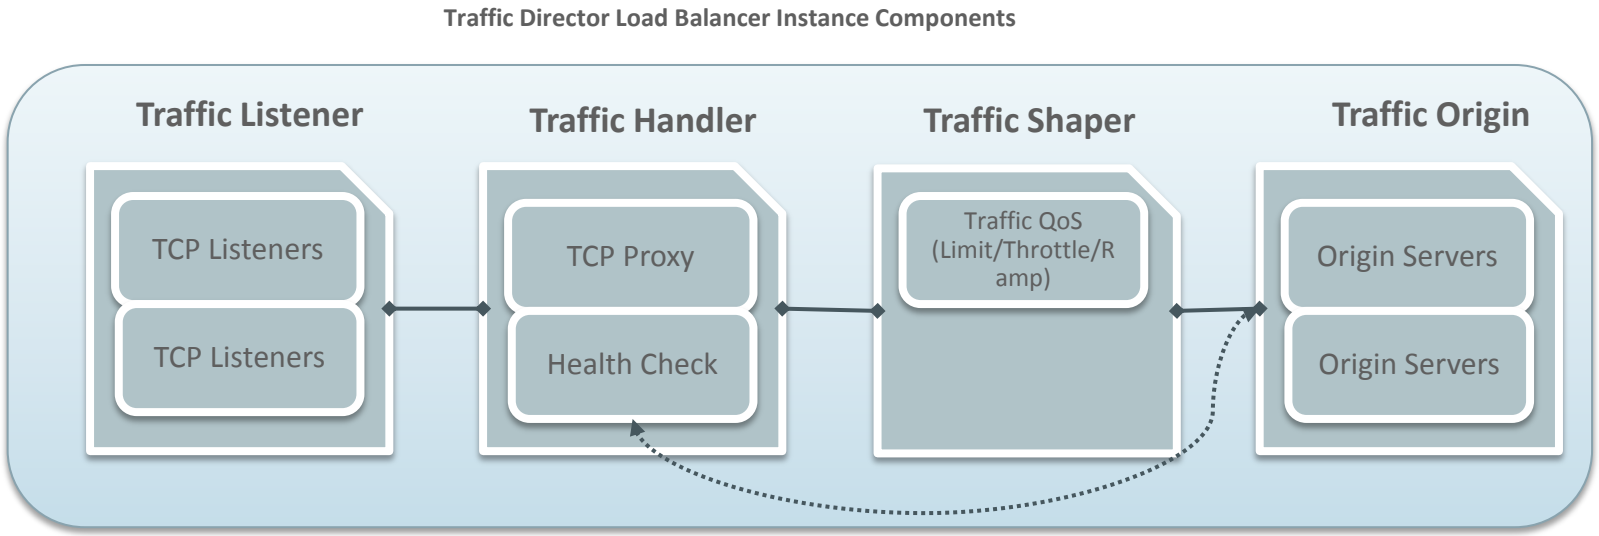

- Traffic Director TCP Load Balancer Server has 4 key components:
  - Traffic Listener - Handles incoming connections (IP/VIP : Port)
  - Traffic Handler - Processes incoming client TCP/IP connections (no data processing)
  - Traffic Shaper - Shapes incoming traffic and can offer QoS for underlying traffic
  - Traffic Origin - Passes traffic to back-end application server where content originates

- Acceptor and Keep Alive threads handle Traffic Listener load
  - SSL Tunneling/Termination
- TCP Thread Pools handle Traffic Handler/Shaper processing load
  - Non-Blocking Thread Pool
  - Origin Server Pool selected based on listener port
- Customizable Health Check
  - TCP/IP Ping
  - External command support for deep health check
- Use Cases
  - LDAP, T3 Initial Connection/Provider (JNDI) LBR

Traffic Director 11g – Creating Load Balancer Policies

Configuration Wizard

- Config Wizard -within administration console- creates the following simplifying the steps to create load balancer policy:
  - New Virtual Server (with default route) with same name as ‘Configuration Name’
  - New HTTP Listener with ‘http-listener-1’ associated with Virtual Server
  - New Server Pool named `origin-server-pool-1` associated with default route for Virtual Server

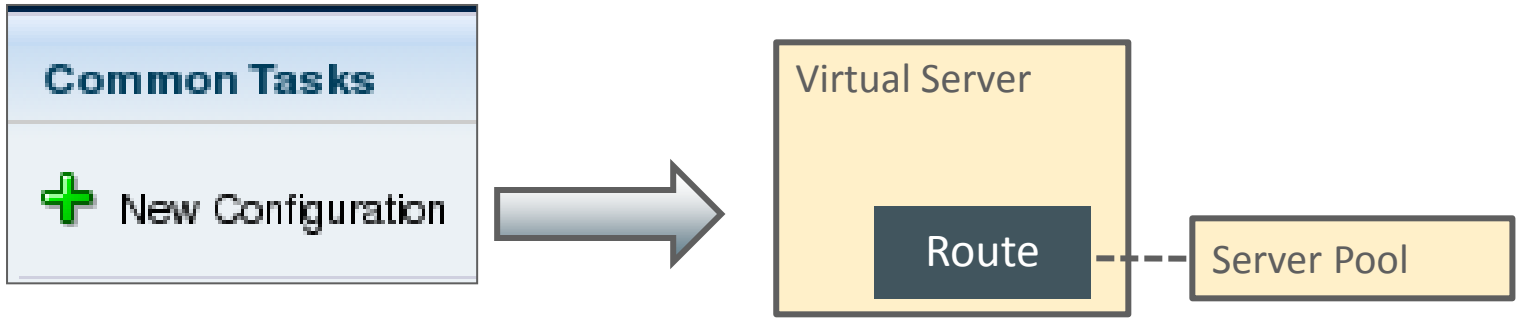

Traffic Director – Creating Load Balancer Policies
Virtual Server Routes

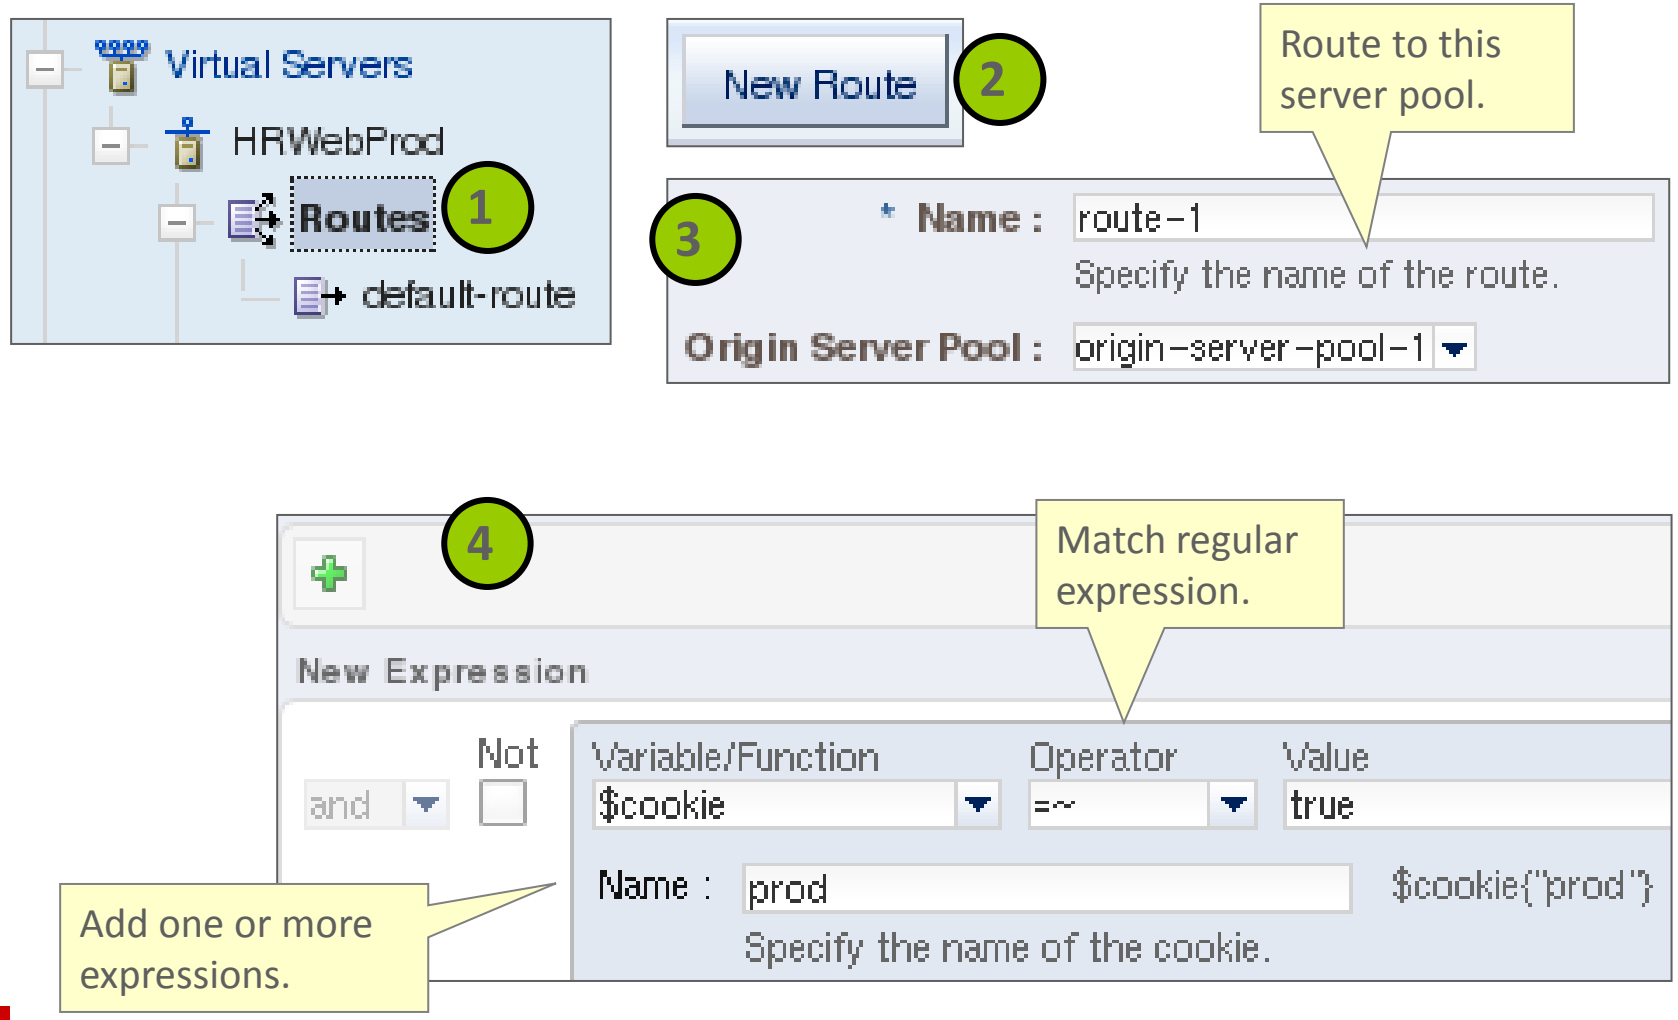

Traffic Director – Creating Load Balancer Policies
Configuring Virtual Server Routes

$browser	Browser Agent
$body	Request Body
$dns	Client DNS
$id	Virtual Server
$ip	IP Address
$keysize	Key Size
$method	Request Method
$protocol	Protocol
$query	Query String
$referer	Referer Header

$url	Request URL
$urlhost	Client Hostname
$chunked	Is Chunked?
$internal	Is Internal Request?
$keep_alive	Is Keep Alive?
$restarted	Is Request Restarted?
$security	Is SSL Used?
$cookie{'Name'}	Cookie Name
$env{'Name'}	Environment Variable
$headers{'Name'}	Request Header

Some of the available variables

Traffic Director – Creating Load Balancer Policies
Administration Console Overview

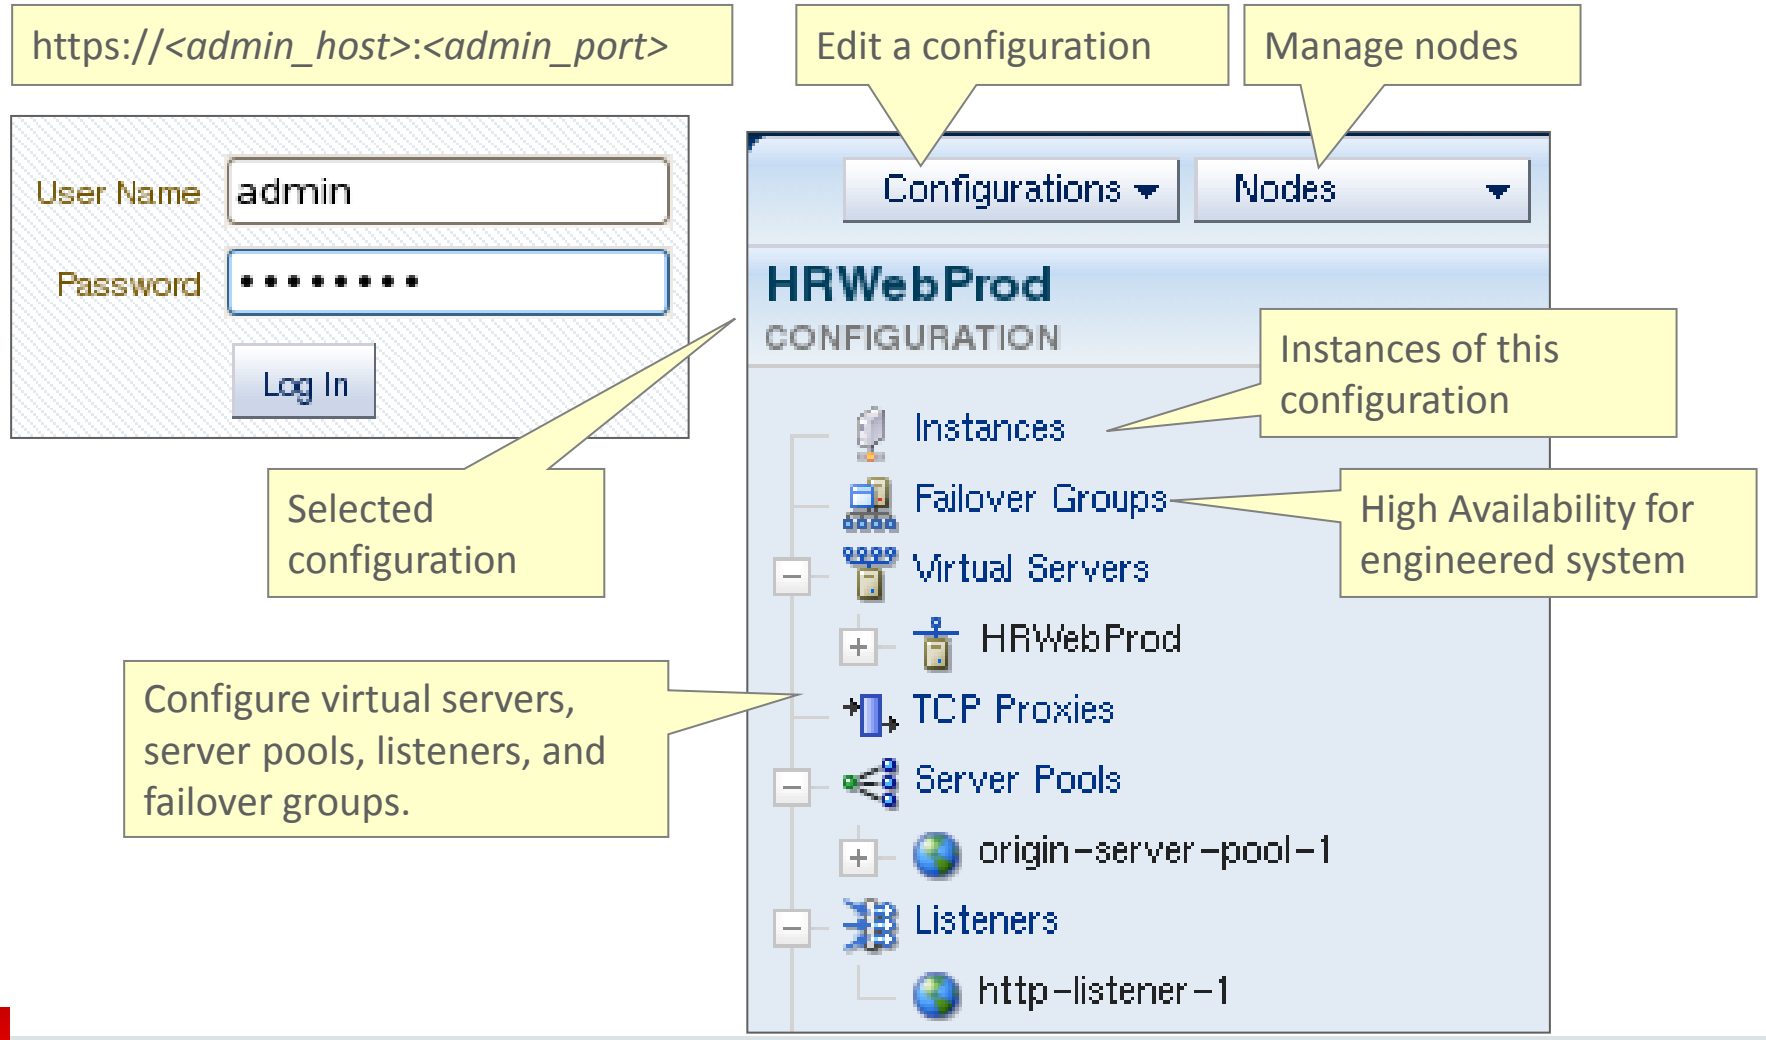

Traffic Director 11g – Creating Load Balancer Policies

Configuring Server Pools

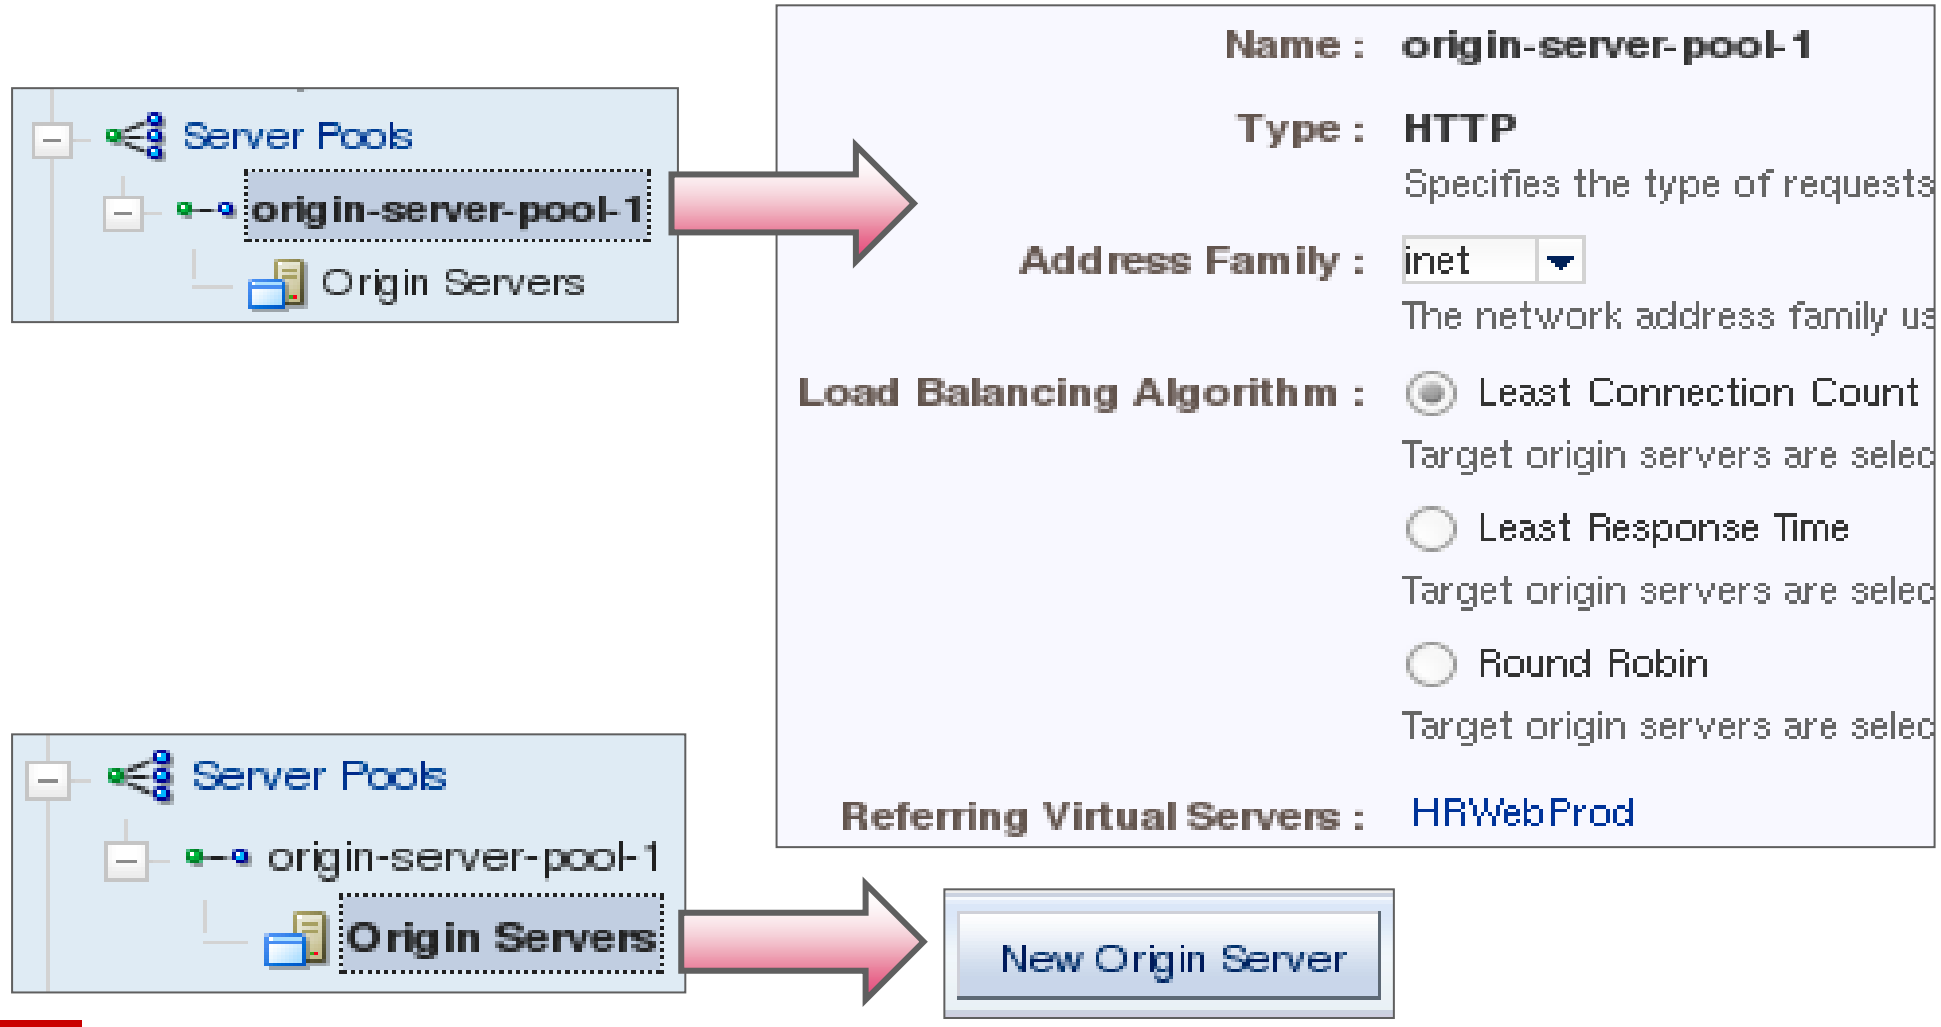

Traffic Director 11g – Creating Load Balancer Policies
WebLogic Cluster Dynamic Discovery

- By default, OTD simply distributes traffic to the current list of origin servers in a pool.
- For WebLogic clusters, OTD can dynamically:
  - Send periodic health checks to the cluster
  - Discover newly added cluster members
  - Update the pool to reflect the latest list of running servers

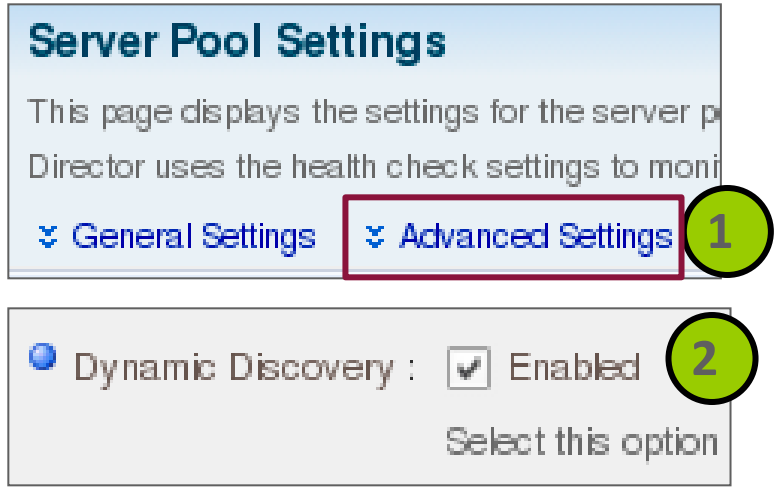

Traffic Director 11g – Creating Load Balancer Policies

Creating Failover Group(s) – Applicable only on engineered systems

- Ensure that all listener addresses are set to * or are the same as the virtual IP.
- Select two running admin nodes (primary and backup).

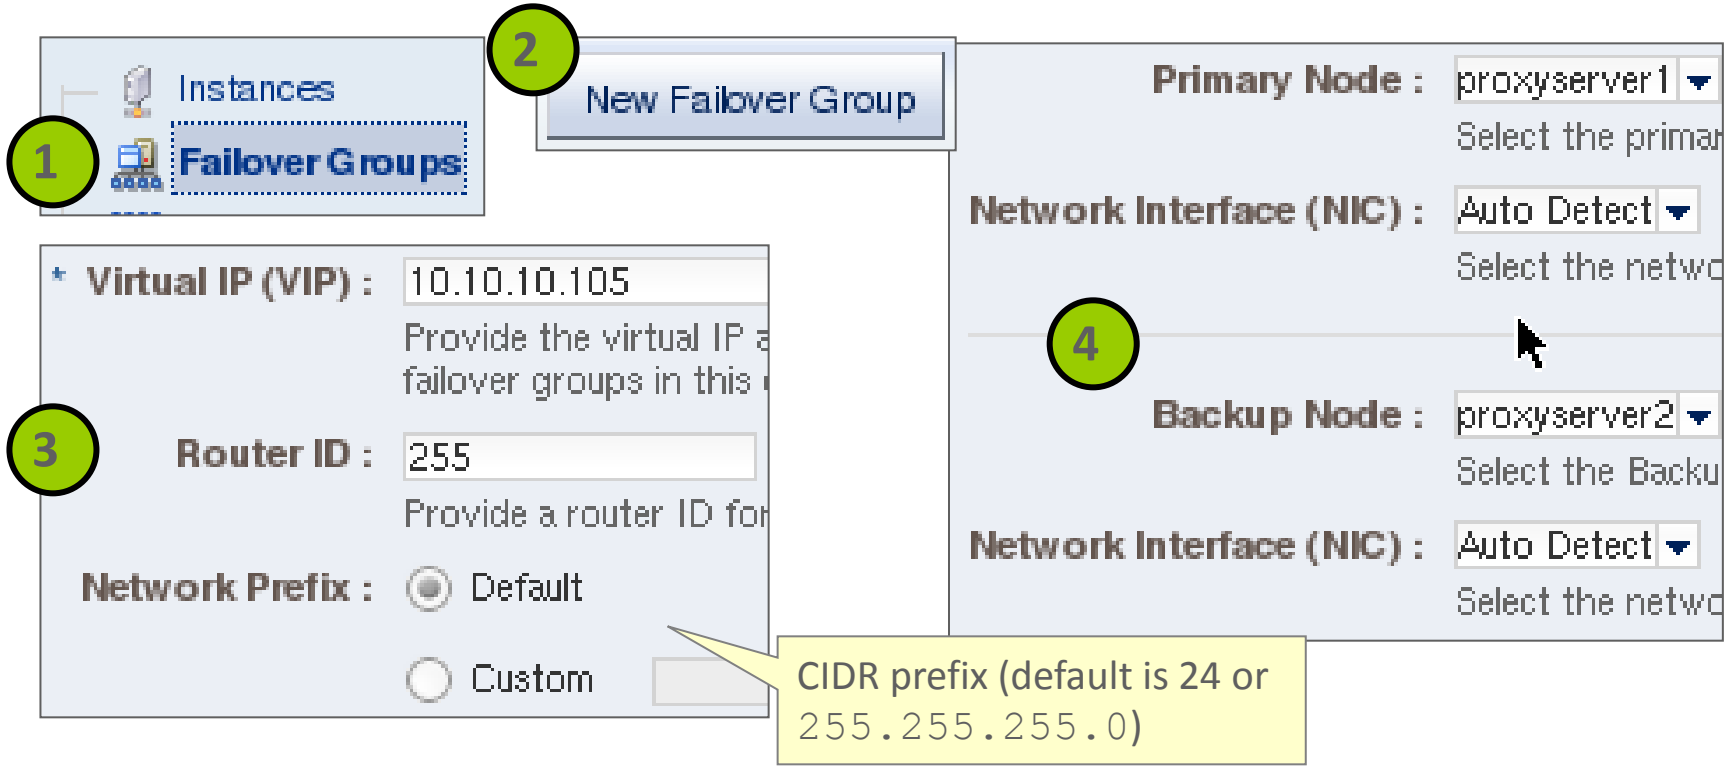

Traffic Director 11g – Administer Server Lifecycle
Starting Failover Group(s) – Applicable only on engineered systems

- If the admin nodes are started as `root`, they will automatically start and stop the `keepalived` process.
- If the admin nodes are not started as `root`, you must manually start `keepalived` as `root`.

Manually start the failover daemon:

```
su -
cd <OTD_HOME>/bin
./tadm start-failover --host=myadminhost --user=admin
   --password-file=admin.pwd --no-prompt --config=myconfig
   --instance-home=/u01/app/otd`
```

Traffic Director 11g – Administer Server Lifecycle
Managing Instances

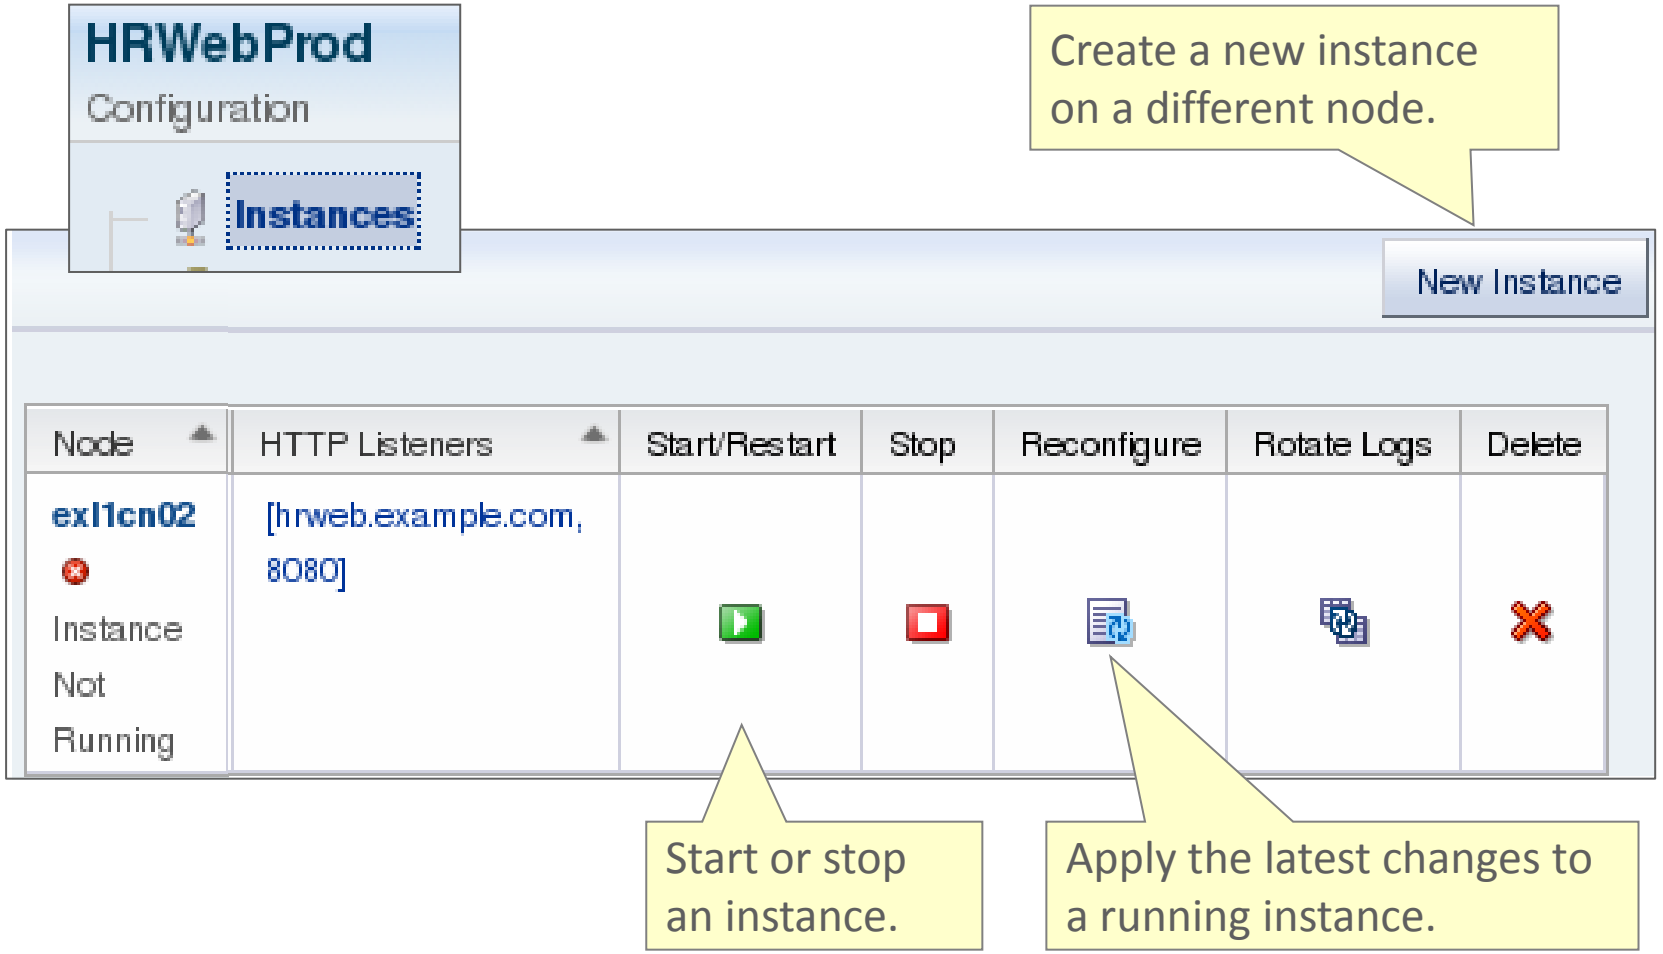

Traffic Director 11g – Administer Server Lifecycle
Log Management

- Each instance has its own log files:
  - *<instance_home>*/*<instance>*/logs/server.log
  - *<instance_home>*/*<instance>*/logs/access.log
- If desired, you can:
  - Configure dedicated log files for a specific virtual server
  - Customize each log's location, output level, or rotation policy

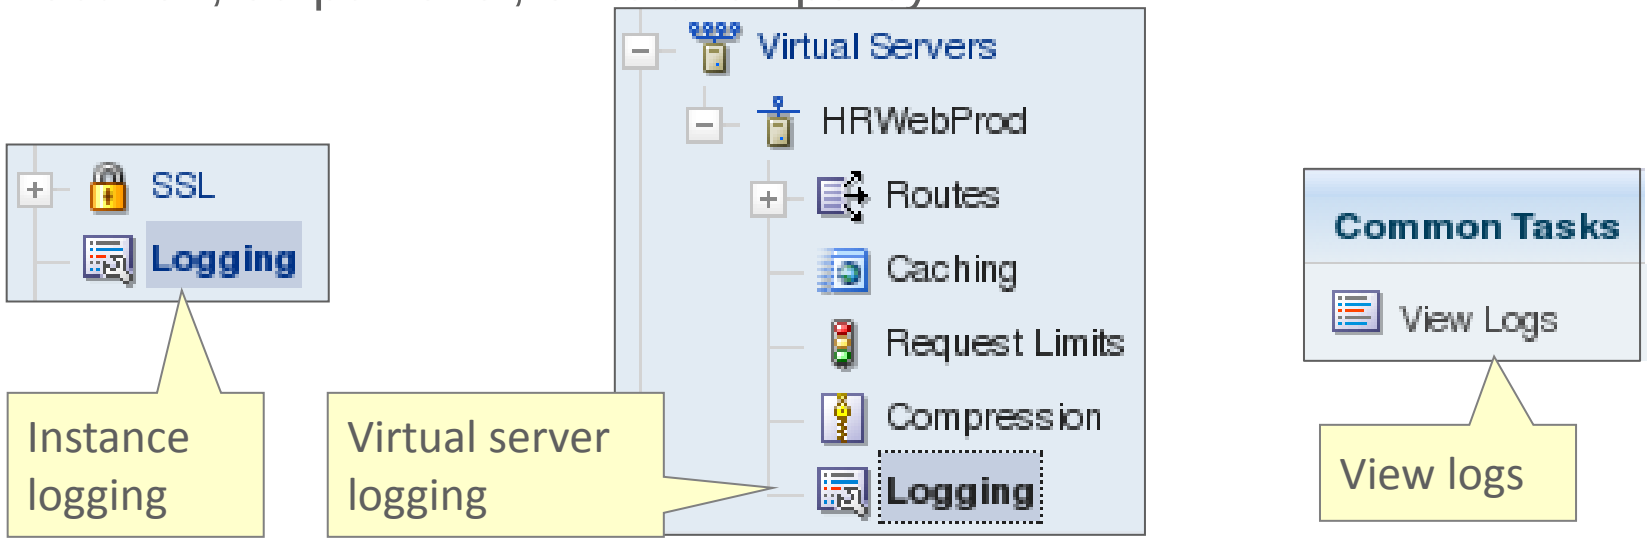

Traffic Director 11g – Process Overview

- Every OTD on a vServer has 2 key components
  - **Admin node** - Special instance that synchronizes with *OTD admin server* and manages *instance configuration and life-cycle*.
    - Admin Server – Special case of Admin Node supporting browser based UI
  - **Server Instance** per OTD configuration deployed on the admin node
- Three ***trafficd*** processes per OTD instance
  - **Watchdog**: Process spawns primordial process; ,manages lifecycle for Primordial and Worker processes (Handles privileged port)
  - **Primordial**; Launches worker processes ; gathers server performance statistics
  - **Worker**: Actual load balancer process serving incoming requests

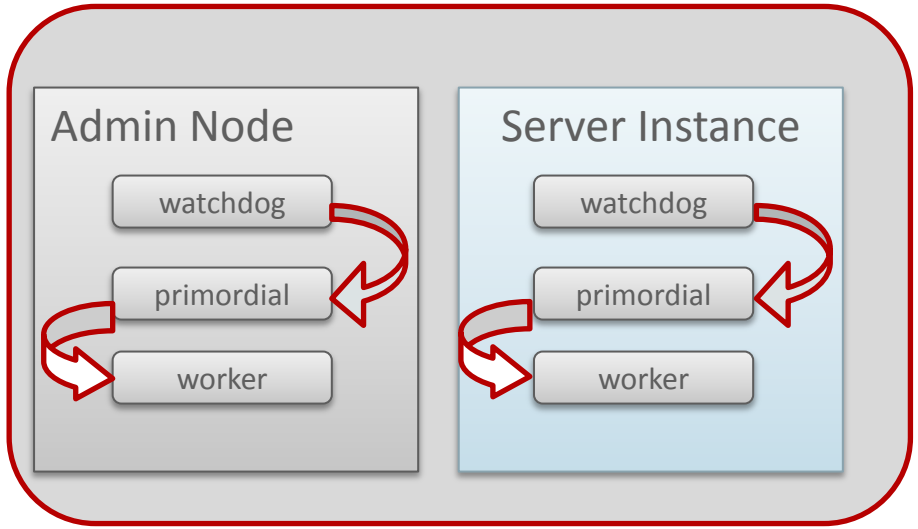


```
[wlsofm@amsooffpe26 ~]$ ps -ef | grep trafficd | grep admin
wlsofm    5665  5100  0 12:24 pts/1    00:00:00 grep trafficd
wlsofm   19259     1  0 May05 ?        00:00:00 trafficd-wdog -d /u01/appl/wlsofm/otd/otdnode2/admin-server/config -r /u01/appl/wlsofm/products/otd
-t /tmp/admin-server-f79dbba8 -u wlsofm
wlsofm   19260 19259  0 May05 ?        00:00:36 trafficd -d /u01/appl/wlsofm/otd/otdnode2/admin-server/config -r /u01/appl/wlsofm/products/otd -t
/tmp/admin-server-f79dbba8 -u wlsofm
wlsofm   19261 19260  0 May05 ?        00:02:52 trafficd -d /u01/appl/wlsofm/otd/otdnode2/admin-server/config -r /u01/appl/wlsofm/products/otd -t
/tmp/admin-server-f79dbba8 -u wlsofm

[wlsofm@amsooffpe26 ~]$ ps -ef | grep trafficd | grep net
wlsofm   23543     1  0 May05 ?        00:00:00 trafficd-wdog -d /u01/appl/wlsofm/otd/otdnode2/net-offort-configuration/config -r
/u01/appl/wlsofm/products/otd -t /tmp/net-offort-configuration-f7683315 -u wlsofm
wlsofm   23544 23543  0 May05 ?        00:00:36 trafficd -d /u01/appl/wlsofm/otd/otdnode2/net-offort-configuration/config -r
/u01/appl/wlsofm/products/otd -t /tmp/net-offort-configuration-f7683315 -u wlsofm
wlsofm   23545 23544  0 May05 ?        00:00:52 trafficd -d /u01/appl/wlsofm/otd/otdnode2/net-offort-configuration/config -r
/u01/appl/wlsofm/products/otd -t /tmp/net-offort-configuration-f7683315 -u wlsofm
```

Traffic Director 11g – High Availability (App vs Node)

Application Level Availability via Watchdog

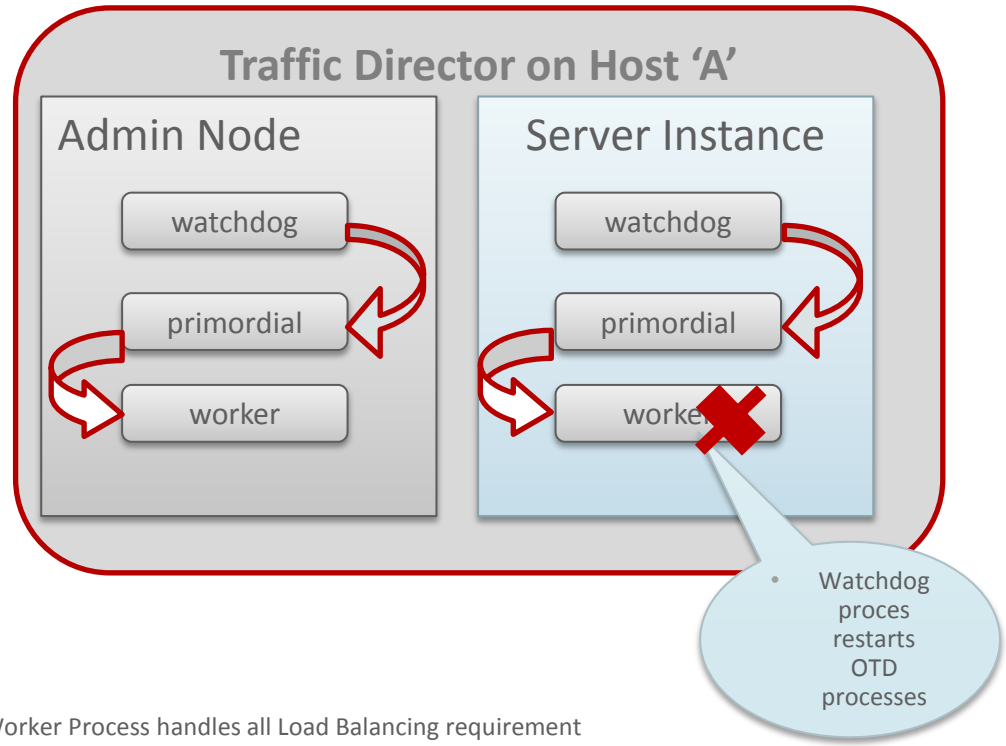

- Worker Process handles all Load Balancing requirement
- Application Availability with Watchdog process (and also OTD EM Agent for EM Cloud Control)
  - Any software crash in the load balancer process (worker process) is handled by watchdog.
  - EM Agent under EM Cloud Control can also monitor Watchdog.

Node Level Availability via Failover Groups

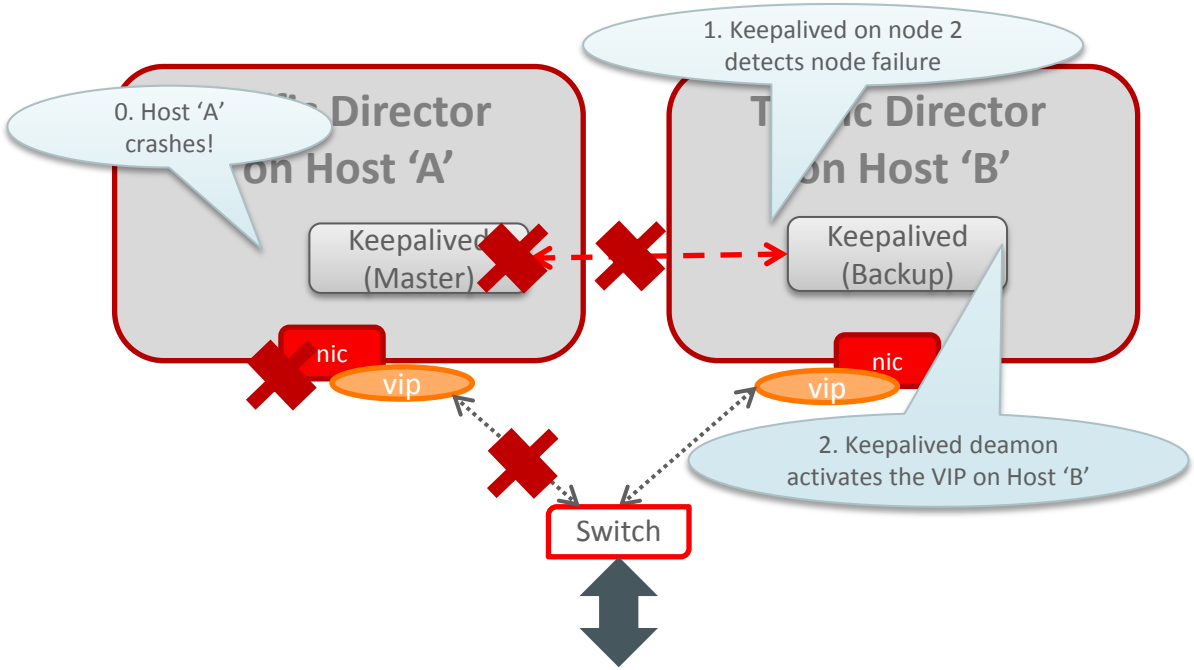

- Avoid Single Point of Failure when Host hosting Traffic Director crashes (OS / CPU / Memory issues!)
- Retain Back-end Application Availability via Failover Groups (VIP on a pair on Admin Nodes)
  - HA Daemon (KeepAlived on Linux) heartbeat determines host crashes with 3 second and takes VIP ownership
  - External Switch / Load Balancer continue to send requests to VIP transparently!

Traffic Director 11g – Monitoring Overview

- Monitoring Statistics provide the following
  - Incoming client traffic (Connection Queue, Requests / Sec, Thread Resources)
  - Outgoing traffic to Origin Server (back-end) Traffic (Connections, Failures, Response Summary - 1xxx, 2xxx, 3xxx, 4xxx, 5xxx)
- In-depth performance stats available accessible via
  - Enterprise Manager Cloud Control
  - SNMP
  - Administration Server CLI ('tadm')
    - get-config-stats / get-virtual-server-stats / get-stats-xml / get-origin-server-stats
- Monitoring Statistics are available in these formats:
  - Text, XML, SNMP
- Enterprise Manager Cloud Control also supports:
  - Traffic Patterns / Alerts

Traffic Director – Monitoring with EM Cloud Control

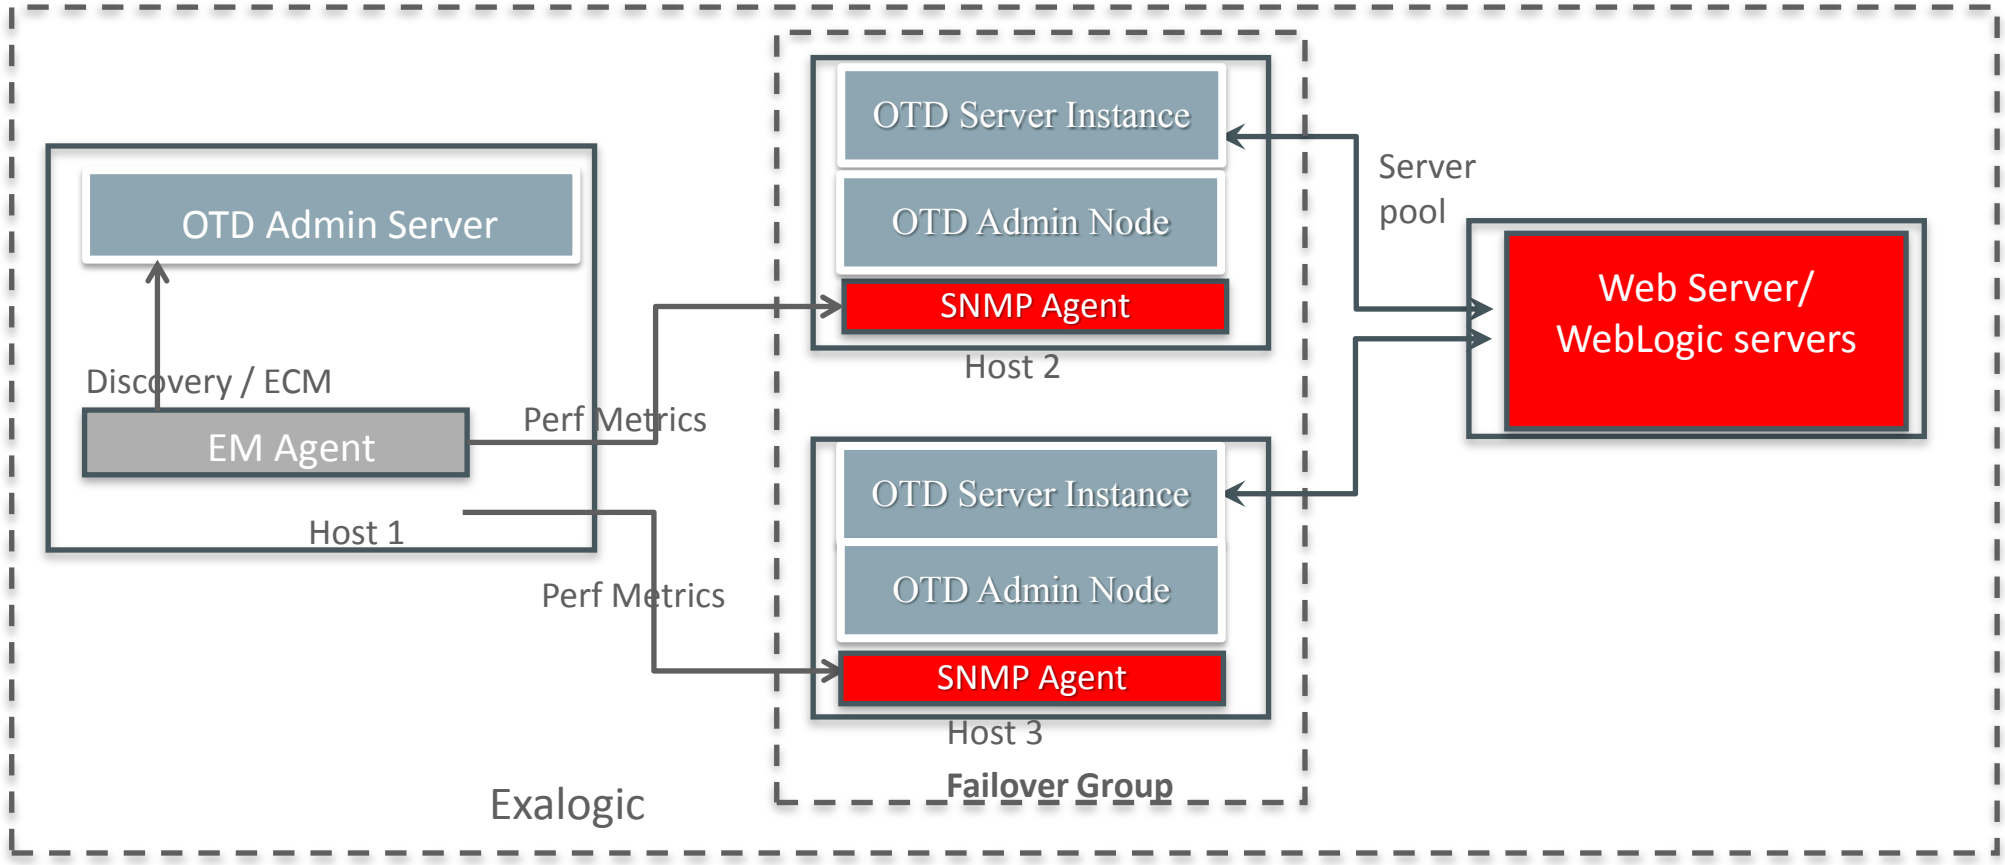

Traffic Director 11g – Technical Deep Dive

- Big Picture View
- Terminology Overview, Topology Discussions
- Feature Deep Dive
- Deployment Best Practices
- Peek at Future Investments
- Q & A

Traffic Director 11g – Production Checklist (1/3)

- Supported System / Configurations (Oracle Support Note: 1910033.1)
  - Uptake latest PatchSet updates before staging / production (Oracle Support Note: 1676256.1)
- Designed for High Availability (No Single Point of Failure)
  - Leveraging Traffic Director built-in cluster capability to ensure No SPoF
    - Traffic Director deployed at least in 2 Nodes (1 Admin Server and 1 Admin Node; Configurations deployed to these 2 nodes)
      - Requires external entity such as HLB or DNS Load Balancer to spray client requests to Traffic Director cluster
    - Admin Server <-> Node communication only via 'host name' (rather than IP Address). Pair these two using the Host name. Allows for IP change with minimal effort
  - Designed to uptake roll-over patches between 2 OTD Admin (Oracle Support Note: 1982144.1)
    - Requires Traffic Director binaries are not shared between OTD Admin Server/Nodes
    - Leverage roll-over patching to apply the same patch on both the Admin Server/Nodes
- Multiple Configurations vs Multiple Virtual Servers
  - Virtual Servers
    - Corresponds to a specific HTTP 'Host' header. Designed to front-end multiple sites / Mass Virtual Site Hosting
  - Configurations
    - Hosts multiple virtual servers within a configuration; Includes separate Load Balancer Service Lifecycle
    - Ideal to front-end multiple environments (Dev, QA, Staging)

Traffic Director 11g – Production Checklist (2/3)

- Configuration Best Practices
  - Enable HTTP response caching when workload / applications use lot of static content (Images, HTML, JavaScript)
  - Enable HTTP response compression when average response size > 8KB
    - Check Traffic Director Load Balancer Service Instance (net-<CONFIG>/config/access.log) to check for the response size
  - Configure Traffic Director (OTD) – Origin Servers (WebLogic) communication via IPoIB (Engineered Systems only)
  - Configure 'Always *Keep-Alive* ' in OTD Route Settings while front-ending SOA, EBS, PeopleSoft applications.
    - OTD 11.1.1.9 enables this by default!
- Process Management - Runtime Process Privileges
  - Does Load Balancing service happen over privileged ports (80/443)?
    - Option 1: Traffic Director Admin Server/Node runs as root; Load Balancer Service still runs as non-root; Lifecycle management only through Admin
    - Option 2: Manage Traffic Director Load Balancer Service lifecycle outside of Administration UI/CLI via 'sudo' .
      - sudo <OTD_INSTANCE_ROOT>/net-<CONFIG>/bin/startserv
      - Actual Load Balancing Service still runs as 'non-root'

Traffic Director 11g – Production Checklist (3/3)

- Monitoring
  - Enable Traffic Director monitoring for monitoring live systems
    - Via CLI (tadm get-perfdump) or SNMP or EM Cloud Control
  - Key Performance Indicators (KPI)
    - Average Queuing Delay: Consistently high value means Increase per process file descriptors and Max HTTP Processing Threads
    - Total Threads vs Active Threads : Active vs Total should not be equal consistently. Increase Max HTTP Processing Threads
- Security
  - Ensure proper patching system in-place for the underlying operating system and Traffic Director software
  - Ensure SSL private keys are handled securely.
    - Public Certificates and Private keys are kept within net-<config>/config
  - Do not enable URI based monitoring on production systems
  - Setup warning to handle Certificate Expiry related issues
- Logging
  - Traffic Director delivers automatic log rotation; Ensure a proper process in-place for log archival.
    - Access Log includes information such as which Client IP and Resource served by which Origin Server; Helpful in troubleshooting.

Traffic Director – Sizing Guidelines (1/2)

Key sizing tips while front-ending enterprise workload

- Sizing Guidelines
  - Understand workload characteristics
    - Avg. Response Size (< 8 KB vs > 16 KB) and Response Time (100–400 milli-seconds vs 1-2 seconds)
    - HTTP vs HTTPS (SSL Termination, SSL Proxy), Is Compression enabled on outgoing HTTP responses?
    - Expected Throughput vs current Network Bandwidth (1 Gb vs 10 Gb)
  - Size OTD (Virtual – vServer/vCPU) based on workload characteristics
    - OTD is I/O intensive. Not memory or CPU intensive.
    - Typical OTD vServer Sizing : 1 or 2 vCPU (depending on throughput requirements) with 8 GB RAM and 8 GB Swap
      - 1 vCPU: 10k HTTP or 5k HTTPS transactions/sec (Average application response time around 100 - 400 milliseconds)
        - Add. 10% overhead (reduced performance) with HTTP response caching, compression enabled!
    - Increase vServer RAM by 1 GB / Swap by 1GB for hosting every additional OTD configurations in a vServer

Traffic Director – Sizing Guidelines (2/2)

Key sizing tips while front-ending enterprise workload

- Tuning Guidelines
  - Load Balancer Service largely runs as Single Process / Multi Threaded.
    - Requires System level tuning to increase number of file descriptors per process
    - Assign at least 32276 file descriptors to OTD process runtime user (ulimit –n)
      - Edit <OTD_INSTANCE_HOME>/net-<config>/bin/startserv and configure ulimit –n value
        - #!/bin/sh
          - ulimit –n 32276
    - Increase per process file descriptor allocation - as necessary - to handle additional client workload.
      - Increase this file descriptor if OTD front-ends large number of origin servers
  - Additionally tune only when necessary!
    - Increase Max HTTP Processing threads when response time >= 1 sec.
      - Configure Max. HTTP Processing threads to 1024 or 2048 when avg. HTTP response time is > 1 sec.
      - Increase Max. HTTP Processing threads to 4096 as you double origin servers
        - Correspondingly also increase the per process file descriptor (say from 65535 to 131072)

Traffic Director 11g – Recommended deployment topology

Increases overall Application Availability while offering maximum throughput

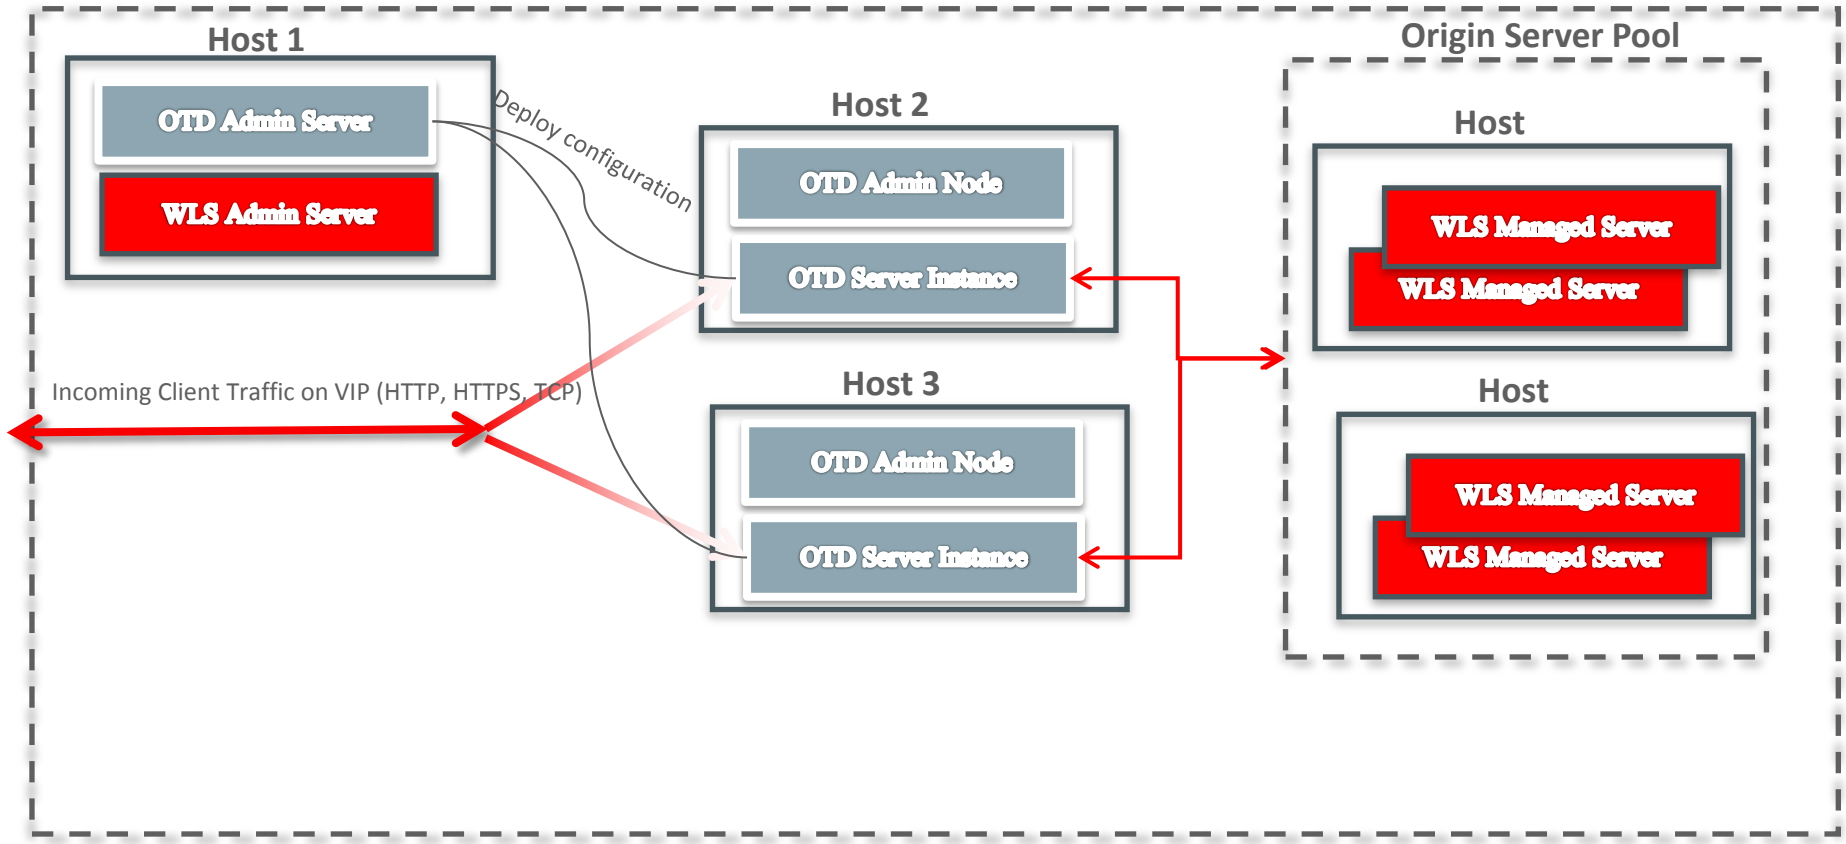

Traffic Director 11g – Topology View (Recap)

Closer Look at a Traffic Director 11g

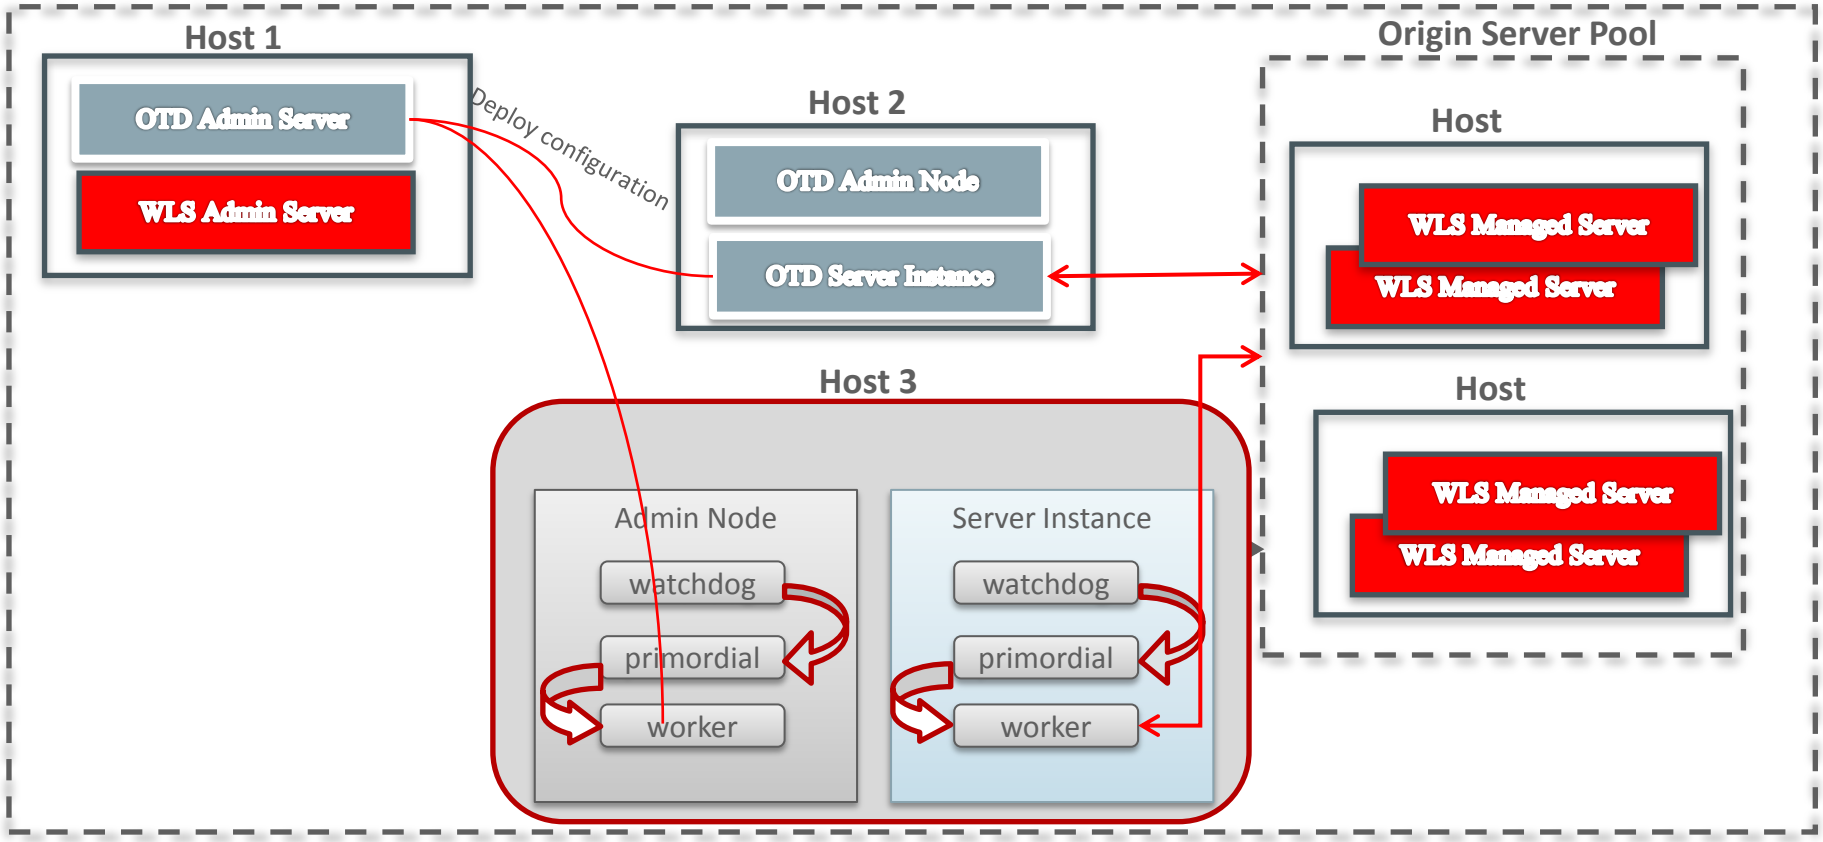

Traffic Director 11g – Technical Deep Dive

- Big Picture View
- Terminology Overview, Topology Discussions
- Feature Deep Dive
- Deployment Best Practices
- Peek at Future Investments
- Q & A

Oracle Traffic Director 12.2.1 – What's New

Focus: Enable WebLogic Multi Tenancy; Oracle ecosystem Integration

- 12.2.1 – GA in Q4 CY2015
  - Integrate with WebLogic Management Framework (New Administration interfaces: WLST, FMWControl)
  - Upgrade toolkit to upgrade from 11g to 12.2.1 - No manual configuration changes.
- Benefits Summary
  - Licensable outside Exalogic (via WebLogic SKU) ; Additional Platform support (Windows / AIX / Linux / Solaris)
  - Support WebLogic 12.2.1 Continuous Availability and Multi Tenancy use cases
    - Seamlessly drain traffic to application server (Zero downtime) ; Elastic scaling of WebLogic dynamic clusters
  - Consistent management interfaces – Fusion Middleware Control and WLST based administration
  - Full-fledged InfiniBand integration – Leverage RDMA for OTD <-> WLS communication (No more sockets!)
  - Full fledged Quality of Service – Traffic Control; Request Limiting; Prioritization; Bandwidth throttling/queuing
  - Serve static content – HTML, Images, JavaScript
  - Bundle WebGate 11gR2 PS3 – Minimizes provisioning overhead

Traffic Director 12c: Recommended deployment in Exalogic

Increases overall Application Availability while offering maximum throughput

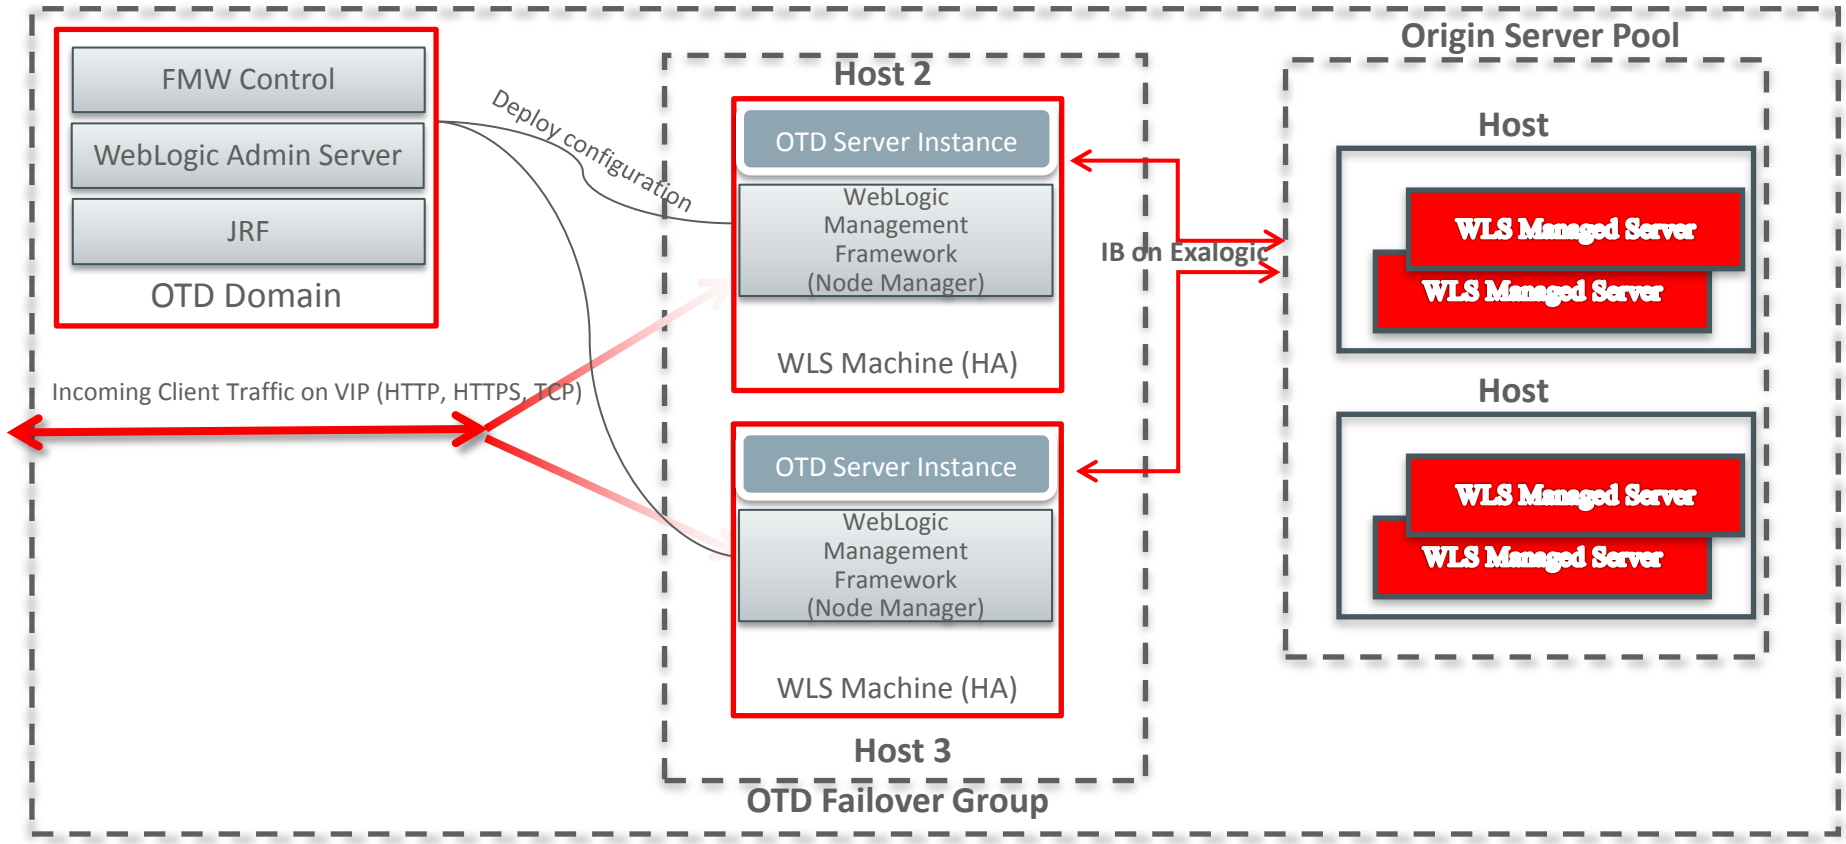

Traffic Director 12c: Recommended deployment in Exalogic

Increases overall Application Availability while offering maximum throughput

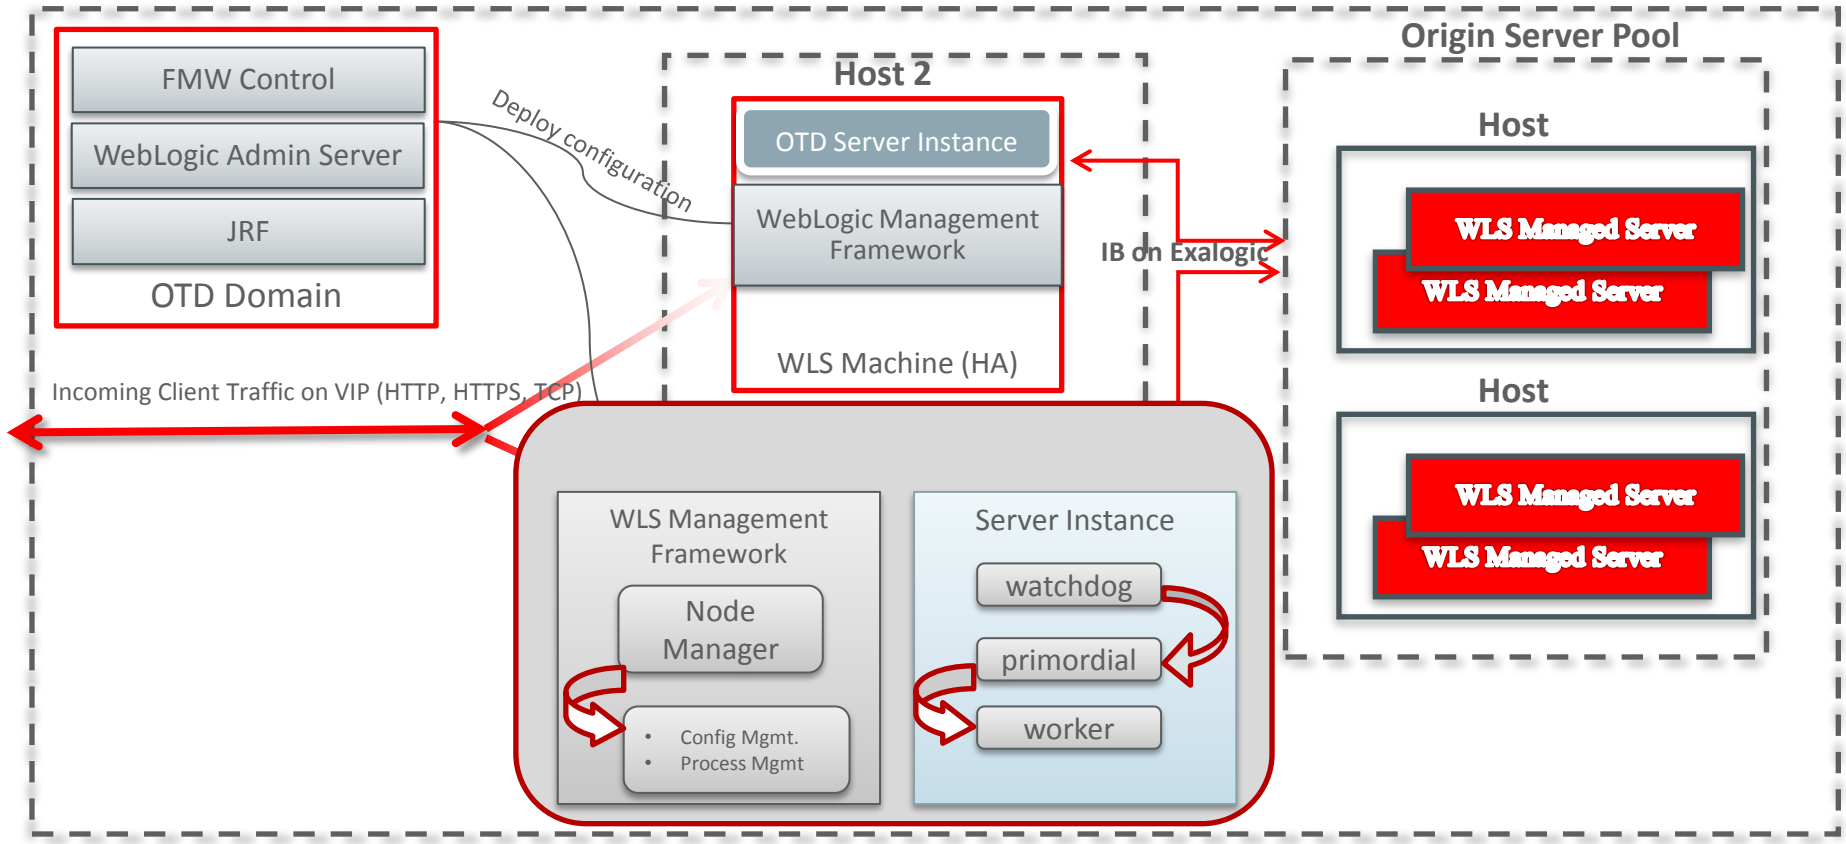

Traffic Director 11g vs 12c – Benefit Analysis

OTD 11g

- Provisioning: Extremely light-weight
  - Layout binaries, Admin Server vs Admin Node
- Administration: Light-weight; Simple UI
  - Browser based UI (JavaScript); No FMWControl
  - New Terminologies Learning curve
  - Lacks: Admin High Availability; Delegated Admin; WLST
    - Automation through *'tadm'* CLI
  - Defaults to self signed certificates
- Core – HTTP / TCP Load Balancing
  - Integration to front-end WebLogic Servers
  - NSS DB stores certificates/keys; No Wallet/KSS

OTD 12c

- Provisioning: Standalone vs Collocated
  - Standalone: No admin post provisioning
  - Collocated: Requires WebLogic Domain + R-JRF
- Administration: Full Oracle Integration!
  - Relevant only on Collocated Provisioning
  - WebLogic based administration
    - WLST; Built-In HA; Delegated Admin support
  - Weblogic Multi Tenancy end-to-end Admin
- Core: Full fledged Quality of Service
  - Bandwidth throttling, Traffic Prioritization
  - Supports Wallets; KSS support in Collocated

Web Tier / Traffic Director – Current Roadmap

CY2013

OTD 11.1.1.7

- TCP Load Balancing
- WebSocket Proxy
- Web Application Firewall
- Solaris Support
- FMW / PeopleSoft Certification

OHS 12.1.2

- WebLogic Management Framework
- Modern SSL Protocol (TLS 1.1+)
- FIPS Compatible

WLS Plug-In 12.1.2

- Apache 2.4 for WLS Plug-In 12.1.2+
- WebSocket Proxy for WLS Plug-In 12.1.2+

CY2014

OTD 11.1.1.9

- Channel Focus (Java Cloud, Access Portal)
- HA Supports on Solaris Zones
- SSL Enhancements (TLS 1.1+/Ciphers)
- BRM Apps Certification

OHS 12.1.3

- WebSocket Proxy (Including. Fallback)
- SSL Offloading Support (SPARC T4+ / Intel)
- Modern Cipher Support

WLS Plug-In 12.1.3

- Apache 2.4 / IIS 8.5 Support
- WebSocket Proxy (Including Fallback)
- Security Enhancements(SSL Offloading/TLS)

CY2015

OTD 12.2.1

- WebLogic Management Framework
- End-to-End WLS Multi Tenancy Integration
- Support for WLS Zero Downtime Patching

OHS 12.2.1

- Based on Apache 2.4
- Ability to front-end WLS Multi Tenancy
- Support for WLS Zero Downtime Patching

WLS Plug-In 12.2.1

- Ability to front-end WLS Multi Tenancy
- Support for WLS Zero Downtime Patching
- Increased Monitoring

Traffic Director – Long Term Roadmap

Key Product Updates & Milestones Anticipated in FY15-16

Fall 2016

Released

OTD 11.1.1.x

- ✓ Load Balancer for Oracle Exalogic and Java Cloud Service
  - ✓ Cluster Administration
  - ✓ Traffic Shaping, QoS
  - ✓ High Availability
  - ✓ Engineered systems focus
  - ✓ OEM Monitoring
- ✓ Load Balancer for Oracle Access Portal, Single Sign-On Suite Plus
  - ✓ Seamless deployment
  - ✓ Available on OEL/RHEL

Q4 CY2015

OTD 12.2.1

- Reverse Proxy focus for WebLogic Multi-Tenancy
  - Enhanced Quality of Service support
  - Integrates Oracle WebGate Plug-In for better SSO support
  - Full Oracle ecosystem integration – FMWControl
  - Multi Platform Support
  - Transparent Proxy
    - Port from iPS 4.0.x
  - Proxy Emulator
    - Port from iPS 4.0.x

Q2 CY2016

OTD 12.2.2

- Load Balancer enhancements for Java Cloud Service
- Reverse Proxy focus for Oracle Identity and Mobile Suite
  - Enhanced Traffic Mgmt. capabilities
  - SSO Plug-In Management

Traffic Director 11g – References

- Deploying OTD within Exalogic – Technical White Paper
- Tuning OTD for High Throughput – White Paper
- OTD with SOA in Exalogic - EDG
- OTD with IDM in Exalogic – EDG
- Disaster Recovery Solution for OTD – White Paper
- OTD – Support Knowledge Base (1626139.2)
- Traffic Director OTN

Technical Deep Dive – Phase 2

Use Cases, Best Practices

Traffic Director – Terminology Recap

- OTD Configuration
  - Metadata information on where and how to accept and process incoming requests
- OTD Instance
  - OTD configuration deployed to an OTD Admin Node (a OTD Server process – handles traffic)
- OTD Administration Server
  - Manage configurations on OTD Administration Nodes
- OTD Administration Node
  - A physical server to which OTD Admin server can translate a configuration to become an instance
- Origin Server
  - Typically - Content (generation) server. Can be another software load balancer!
  - A server in the back end to which OTD proxies HTTP(s) traffic that it receives from the client
- Origin Server Pool
  - Collection of origin servers that hosts same service that you can load balance with OTD
  - Availability monitored by OTD's periodic health checks (over HTTP or TCP)
- Virtual IP (VIP) and Failover Group – Only on Oracle platform - Exalogic, SuperCluster, ODA, Java Cloud Service
  - A pair of OTD instances working in tandem (active/passive) to provide VIP and IP failover based on VRRP

Traffic Director – Configuration(s) vs Virtual Servers(s)

Configuration

- Configuration deploys to LB service with:
  - configuration files under net-<config> directory
  - 3 processes (trafficd-wdog/trafficd/trafficd)
  - Deploys to one or more Administration Node(s)
- Offers complete isolation
  - Separate Config and Server Lifecycle Mgmt.
- Multiple configurations (and LB service) cannot listen at same IP:Port (share Listeners)
- Ideal to front-end multiple environments (Dev, QA, Staging etc.)
- Heavy weight.
  - Needed only when separate lifecycle Mgmt. is critical

Virtual Server(s)

- Ideal to front-end multiple Host based end points
  - Every Virtual Server handles traffic based on 'Host' information in HTTP header
- Virtual Server(s) are included within Configuration
  - Every configuration has one or more Virtual Severs
  - Virtual Server uses separate Listener or share same IP:Port
- Light-weight
  - Ideal to front-end Mass Virtual Host Load Balancing Use Case
  - Ideal to front-end Multiple Hostnames (app1.partner, app2.partner) within the same environment (say Staging)
  - Able to separately disable Virtual Server to stop processing

Traffic Director – Mass Virtual Hosting
Deployment View

- Create HTTP Load Balancer configuration in OTD with:
  - HTTP Listener (accept incoming requests)
  - Origin Server Pools (content origination end point)
- Now, configure one or more Virtual Servers (VS) where every VS
  - Listens to unique HTTP listener or share an existing listener
  - Handles unique 'Host' header or multiple 'Host' headers
- Configuring Routing Rules for every VS to shape incoming traffic.
- Configure DNS alias so OTD VIP receives request for one or more DNS names (partners*.oracle.com)

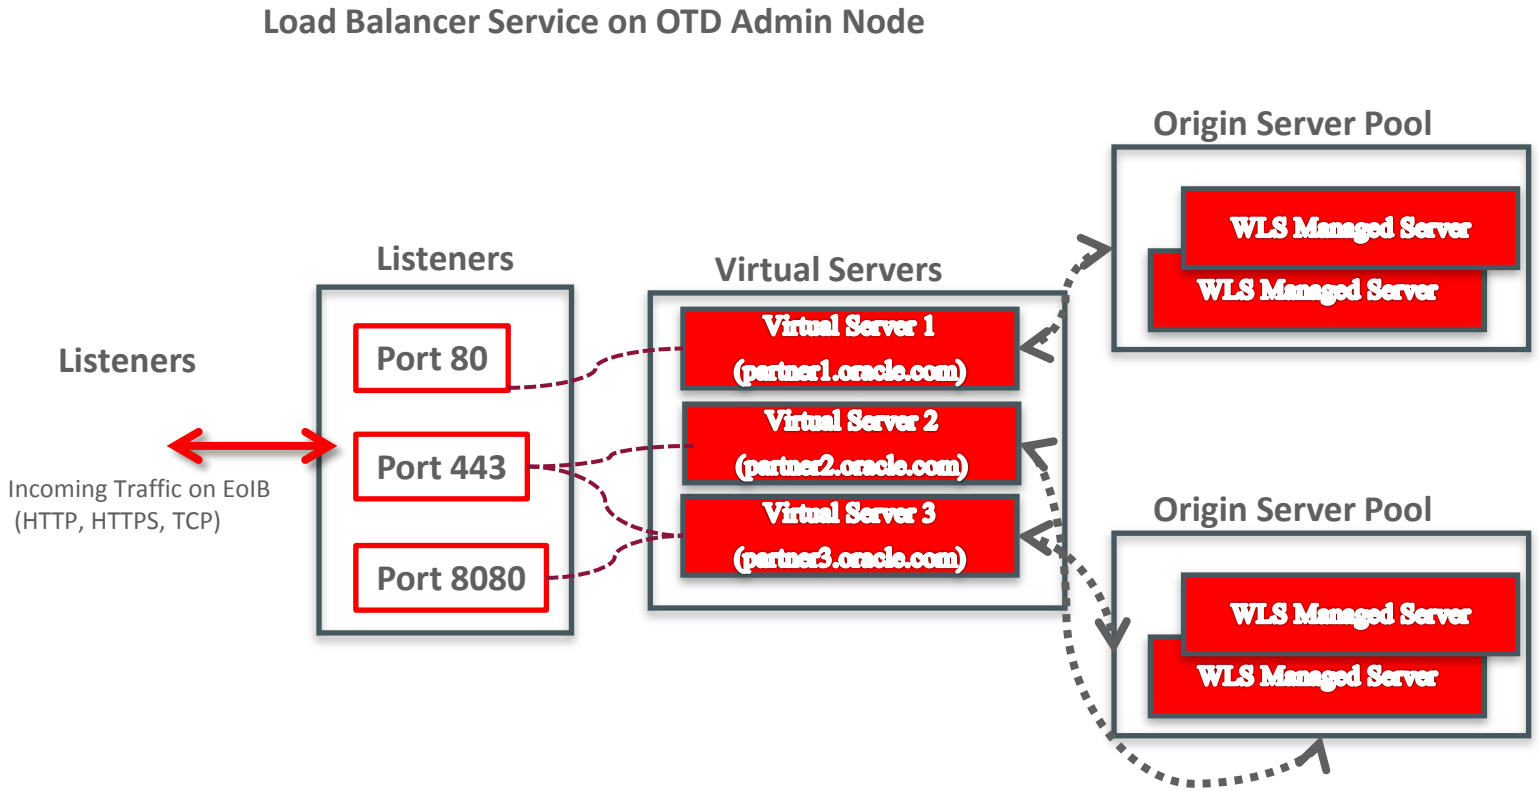

Traffic Director – Scaling Up! (1/2)

Use Case – Scaling up to 30k simultaneous HTTP(s) transactions / sec.

- You cannot optimize what you do not measure!
- Measure system resources
  - Measure OS file descriptor(s) allocated to Traffic Director runtime process (ulimit -n)
  - Allocate additional file descriptor(s) to OTD Server Instance to at least 1013678
  - Allocate at least 4 vCPU with 8 GB memory and 8 GB swap space to OTD vServer
    - With external DNS load balancing, OTD HA setup becomes Active – Active allowing each OTD vServers to have only 2vCPU
- Actively monitor OTD statistics – Necessary to measure OTD performance
  - Key resources to measure within OTD monitoring report
    - Avg. Connection Queuing, Active HTTP Sessions, Persistent connections with Origin Server(s), Origin Server Connection Timeouts
  - Increase Max. HTTP Processing Threads when monitoring data highlights *Queuing Delay*, *Active = Total Sessions*
    - Configure Max. HTTP Processing Threads to as high as 8192 depending on how many origin servers and avg. HTTP response time
    - Explore OTD Multi Process setup for high scalability with slow HTTP response applications.

Traffic Director – Scaling Up! (2/2)

Use Case – Scaling up > 50k simultaneous HTTP(s) transactions / sec.

- Optimize system resources
  - Increase allocated file descriptors to OTD Server Instance process to 2097152
  - Allocate at least 4-6 vCPU with 8 GB memory and 8 GB swap space to OTD vServer
    - With external DNS load balancing, OTD HA setup becomes Active – Active allowing each OTD vServers to have only 3 vCPU
  - Increase Ephemeral Ports within OTD vServer
    - Add this entry within /etc/sysctl.conf
      - *net.ipv4.ip_local_port_range = 1024 65535*
    - Run */sbin/sysctl –p*
- Optimize OTD resources
  - Key resources to measure within OTD monitoring report
    - Avg. Connection Queuing, Active HTTP Sessions, Persistent connections with Origin Server(s), Origin Server Connection Timeouts
  - Increase Max. HTTP Processing Threads when monitoring data highlights *Queuing Delay*, *Active = Total Sessions*
    - Configure Max. HTTP Processing Threads to as high as 16192 depending on how many origin servers and avg. HTTP response time
    - Explore OTD Multi Process setup for high scalability with slow HTTP response applications.