



## Multi-slice scalability improvements in HPSA 7.8

Last Update Date: 09 June 2009

### Overview

The scalability of HP Server Automation was significantly improved in version 7.5 with the ability to deploy multiple instances of key SA components within a single core. This allows multiple instances of the twist and way to operate in parallel. Further scalability improvements within 7.8 allow multiple instances of the word component to be distributed across multiple servers within a core. This paper provides a case study on improved scalability for SA 7.8 in the context of a common SA job type, Application Configuration.

### Executive Summary


The stated test case was exercised against a 7.8 GA Server Automation core. The test was executed against configurations which had between 1 and 3 slices participating in the core. The below table includes the job throughput characteristics for each test configuration:

	Throughput (servers/minute)
# SA Core Slices	7.8 GA
1	96
2	154
3	254

**Table 1: Peak Throughput for Application Configuration, over number of Slice Servers**

For these tests, throughput is calculated by dividing the overall job completion time by the number of managed servers in the AppConfig job submitted for processing.

The following table clearly outlines the improved scale-out characteristics with the addition of multiple slices to the core.





## Test Case Description

The implemented application configuration policy test case:

1. Pushes 3 values down to an /etc/hosts-like config file; the target file has 100 pre-existing entries
1. Contains 1 configuration template
  - a. Standard /etc/hosts configuration template
  - b. We preserve formats and values
  - c. We have 3 default values to get pushed to the managed server.
2. The test case will be run via the UAPI through pytwist.

## Scalability Notes

### Application Configuration Scalability within SA Core

For the Application Configuration Job submissions, SA Core will automatically distribute workload across all instances of the SA "Slice" server within the Core. In this way, AppConfig throughout of large job submissions can benefit from the horizontal scalability of the Slice server. Overall throughput increases as the number of Slice servers is increased.

The following series of graphs plot the execution time periods of the operations on each Managed Server, for the same Job submission of a 100-server workload, as the number of Slice servers is increased from 1 to 3 in the 7.80 core. The X-axis plots the time required for each Managed Server and for the entire Job. Each Wayscript thread corresponds to a thread within the core running the AppConfig update of an individual Managed Server.

Overall Job time is decreased as workload is distributed across multiple Slices.




Figure 1: 100 Server Job, 1 Slice Core

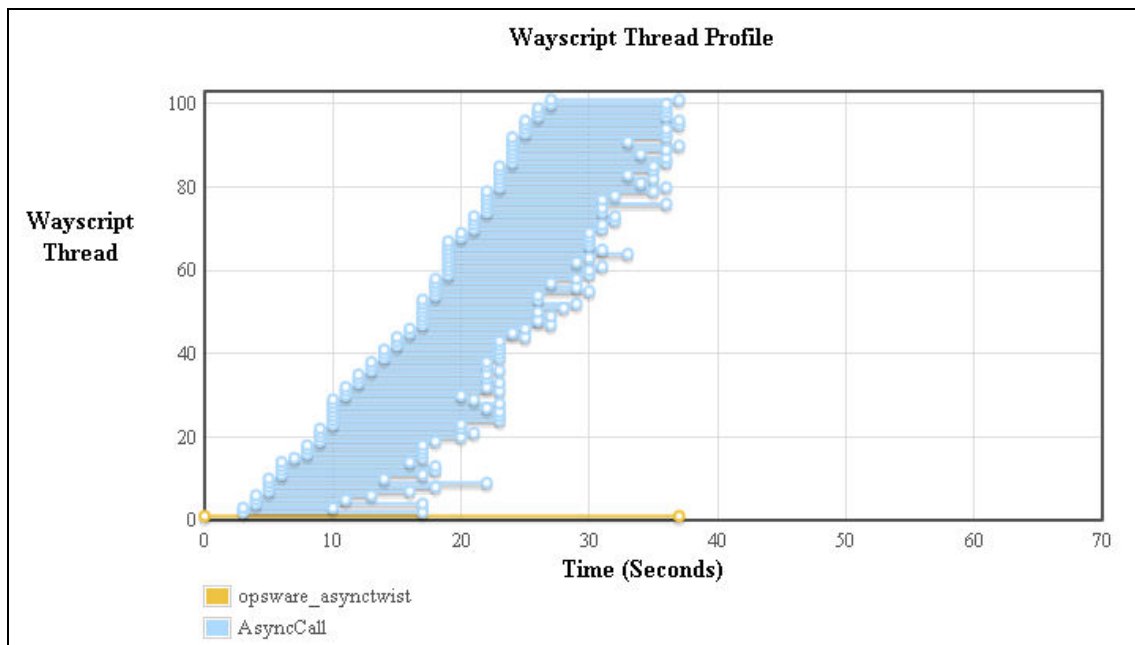


Figure 2: 100 Server Job, 2 Slice Core




Figure 3: 100 Server Job, 3 Slice Core

## Conclusions

The total throughput of Application Configuration increases as the number of Slice servers is increased, showing horizontal scale-out of the Slice server in SA Core.



## Hardware Configuration

### SA Core Servers

Server Role	"Truth" Database
Hardware Specs	Local Disk: 2x 72GB 10K SA RAID-1 (Linux boot) Local Array: MSA50 12x 72GB 10K SAA, Ultra-SCSI connect Memory: 32GB PC5300 OS: RHEL AS 4 64-bit CPU: 2x Dual-core 2.66 GHz. Intel Xeon 5150 Model: HP Proliant DL360G5
Network Config	Network: 1 Gbps LAN, quiet VLAN
Software Specs	Oracle 10.2.0.2.0 Standard Edition
Additional Notes	
Server Role	"Infrastructure" services Data access engine (Spin - primary) Master Gateway Media Repository (NFS,SMB)
Hardware Specs	Local Disk: 2x 72GB 10K SA RAID-1 (Linux boot, SA installed) SAN Attach: 2Gbps dual path FC, MSA1500 Array (Media store) Memory: 16GB PC2-5300 OS: RHEL AS 4 64-bit CPU: 2x Quad-Core 2.66 GHz Intel Xeon 5355 Model: HP BL460cG1
Network Config	Network: 1 Gbps LAN, quiet VLAN
Software Specs	SA 7.8 Build 7.0.2679.0
Additional Notes	
Server Role	"Slice" scalable services. Scaled-out from 1-3 instances as noted. Distributed Command Engine (Way) Spin – secondary Web service API (Twist) Opsware Global File system (Hub) Spoke Word Build Manager Gateways
Hardware Specs	Local Disk: 2x 72GB 10K SA RAID-1 (Linux boot, SA installed) SAN Attach: none Memory: 16GB PC2-5300 OS: RHEL AS 4 64-bit CPU: 2x Quad-Core 2.66 GHz Intel Xeon 5355 Model: HP BL460cG1
Network Config	Network: 1 Gbps LAN, quiet VLAN
Software Specs	SA 7.8 Build 7.0.2679.0
Additional Notes	

**Table 2: SA Core System Configuration**

**Managed server farm characteristics**

Server Types	<ul style="list-style-type: none"> <li>▪ Blade servers hosting linux VMware Virtual Machines</li> <li>▪ Each VM hosting "Simbots"</li> <li>▪ "Simbots" are chroot-ed environments running under the Linux OS. They are virtual representations of managed servers which can share a common physical host. They are used to simulate large numbers of managed servers in the lab environment.</li> </ul>
Hardware Specs	Local Disk: 2x 72GB 10K SA RAID-1 (ESX boot) SAN Attach: 4Gbps dual path FC, MSA2012 Array (VM images) Memory: 32GB PC2-5300 OS: VMware ESX Server 3.5.2 CPU: 2x Quad-Core 2.66 GHz Intel Xeon 5355 Model: HP BL460cG1
Managed Server Specs	SA Scalebots: Linux X86 2.6.9-34 4x Virtual CPUs 8 GB memory 40 GB disk
Configuration Details	<ul style="list-style-type: none"> <li>▪ Pools of simbots are distributed across up to 6 physical Blade Servers for each test.</li> <li>▪ All device groups will be distributed as evenly as possible across the VMware VM servers during the runs.</li> </ul>
Network	Network: 1 Gbps LAN, HP ProCurve Switch, quiet VLAN
Additional Notes	

**Table 3: Test System Managed Server Configuration**