



Primary Response 3.1 Server Agent Frequently Asked Questions

GENERAL

What is Primary Response?

Primary Response, industry leading intrusion prevention software (IPS), automatically detects, classifies and responds to complex threats, accelerating time to protection and enabling IT to deliver business continuity without compromising visibility and control.

What are the key enterprise challenges that Primary Response addresses?

Enterprise challenges include:

- Need for comprehensive real-time protection for PCs and servers across heterogeneous environments
- Guard critical information assets and data
- Ensure regulatory compliance
- Mitigate management pain and costs
- Handle mobile and occasionally-connected end-users

What makes Primary Response unique from any other application security solution?

Unlike other technologies that offer only rules or signatures, Primary Response offers a combination of IPS techniques, used simultaneously, to prevent the most complex and unpredictable attacks using (1) out-of-the-box knowledge-based system to protect end-users, applications and services, (2) Active Malware Defense Technology (Active MDT) which are behavioral heuristics that can detect and prevent malicious software (malware) activity and (3) adaptive profiling that learns application file path behaviors and responds to anomaly-based threats with Sana Adaptive Profiling Technology (SanAPT). Taking a preemptive protection approach, Primary Response provides 'always on' protection for both PCs and servers and requires no new signature updates to prevent the latest, emerging threats, which enables IT to deliver business value across the enterprise.

How does Primary Response work?

Primary Response offers a combination of innovative IPS techniques, used simultaneously, to prevent the most complex and unpredictable attacks using (1) out-of-the-box knowledge-based system to protect end users, applications and systems, (2) Active Malware Defense Technology (Active MDT) that can detect and prevent suspicious malware activity such as Trojans, keyloggers, silent backdoors and rootkits based on a behavioral heuristic approach and (3) Sana's Adaptive Profiling Technology (SanAPT) that provides instant protection for memory-based attacks and learns normal application file path behaviors and responds to anomaly-based threats.

Primary Response Agents are centrally managed from a web-based Management Console using an encrypted and authenticated protocol.

What are the key benefits of Primary Response?

Primary Response offers unmatched capabilities around protection, enforcement, visibility and integration; this directly translates into key benefits for both your enterprise and IT administrators. For some of Primary Response's key benefits, see <http://www.sanasecurity.com/products/pr/features.php>.

How can I see a demo of Primary Response?

Contact Sana Security sales department at 866-900-SANA for a one-on-one demonstration.

How can I get an evaluation of Primary Response?

Contact Sana Security sales department at 866-900-SANA to receive a time-limited evaluation copy of Primary Response.

Who should I contact at Sana for more information?

Contact Sana Security sales department at 866-900-SANA.

Defense

Why is Primary Response better than competitive products that also claim behavior-based intrusion prevention?

Security products such as McAfee Enterecept and Cisco Security Agent use only knowledge-based policies (rules or signatures) to predefine what an application can and cannot do, and what constitutes an attack. These companies claim that rules and signatures can adequately model normal and anomalous application behavior. While a rules-based approach can be effective in firewalls and other network-focused security products, Sana believes that the other companies' approach is fundamentally flawed when it comes to application security. Application behavior is significantly more complex and cannot be accurately defined with simple rules and signatures.

Because rules and signatures are not granular enough to correctly define normal and anomalous application behavior, they force the administrator to choose between an effective security policy and a low number of false alarms. Locking down an application with too many rules will prevent it from meeting its core business objectives and increasingly becomes hard to maintain. On the other hand, modifying or turning off rules to enable the application to run unfettered, or to lower false alarm rates, enables potential vulnerability exploits. Moreover, because IT environments are diverse and constantly changing, rule and signature modifications are necessary for each server with every system change. This includes operating system and application upgrades, security patches, and configuration changes, making the rules-based solutions not only inaccurate, but also impractical to deploy across the enterprise on production systems.

What are code injection or buffer overflow attacks?

When an attacker overflows the buffer with programmatic instructions, the attack is a **code injection**. The injected code attempts to take over the machine so the attacker can access private data or use the machine to attack still more machines. Primary Response provides protection against the most common forms (85-95%) of code injection attacks, such as those used by Blaster, Slammer, Sasser and Code Red.

How does Primary Response prevent code injection or buffer overflow types of attacks?

Primary Response knows what memory has been allocated to legitimate process. If code begins to run from another part of memory, Primary Response takes specific steps. A code injection attack can take one of two forms:

- In the first type, all malicious code is injected into the buffer.
- In the second type, only a very small amount of malicious code is injected into the buffer. This small piece of code then activates normally benign code for a malicious purpose.

Primary Response observes when system calls occur in invalid memory spaces. In the first type of attack, the code must execute a fairly high number of system calls to accomplish anything. If the number of system calls executed matches the Minimum System Calls filter set in Primary Response for that application, then Primary Response generates an alert. In addition, Primary Response will block additional system calls if it has been configured to do so. In the second type

of attack, the code launches an external process. Primary Response will block this behavior regardless of the number of system calls used by the attacker.

I already have a firewall and a Network Intrusion Detection system at my perimeter. Do I still need Primary Response to protect my servers and PCs?

Yes. Perimeter defenses, such as firewalls and NIDS, provide an important security measure for computer networks. However, they do not provide a sufficient level of protection for PC and server applications. Because many applications communicate with each other and with end users over the Internet, application-level attacks will often penetrate a perimeter via a legitimate access point. Moreover, firewalls and NIDS are unable to inspect encrypted (SSL) traffic, which is not decrypted until it reaches the host. Attacks that successfully navigate the perimeter security layers are often targeting specific application vulnerabilities and are easily detected by Primary Response. According to Gartner Group, 75% of successful attacks exploit application vulnerabilities. Most security experts recommend a layered approach to information security. Primary Response is an application security solution that complements NIDS and other network-layer products, and protects both PC and server-based applications more effectively and accurately.

What is the difference between Primary Response and an application firewall?

Application firewalls are designed to protect web traffic by inspecting the incoming and outgoing HTTP/HTTPS packets against a set of predefined web server-centric filters. Because they are not capable of monitoring actual (web) application execution processes, application firewalls require difficult manual fine tuning to protect highly customized web servers. In addition, application firewalls are susceptible to spoofing, evasion and Distributed Denial of Service (DDoS) attacks.

We have invested heavily in our signature-based solution. Why should we spend more for another product?

Signature-based approaches require previous knowledge of an attack and require a developed signature for a specific type of attack. Products that use a signature-based approach must have up-to-date signatures installed in order for this type of system to work. Of greater importance, a signature-based approach, by its very nature, will fail at preventing zero-day, unknown attacks.

The real cost of signature-based solution is measured in the amount of time required to (1) install and configure signature updates in initial deployment, (2) manage the ongoing signature updates and general change management in the IT environment, and (3) loss of business continuity, productivity, and information data resulting from a successful penetration from an unknown attack, which signature-based systems neither detect nor prevent.

I have heard about McAfee Enterscept and Cisco Security Agent (formerly Okena). Are these products similar to Sana's Primary Response?

No. Primary Response is uniquely different from either McAfee's or Cisco's IPS solution where those legacy solutions provide only signatures or rules respectively. Traditional security systems fall short because they scan data or rely heavily on signatures for detecting malware and other types of attacks, leaving the enterprise potentially vulnerable. The key to solving the enterprise security challenges is to monitor behavior (with behavioral heuristics), and to use a combination of techniques to ensure that any malicious attack can be prevented. Using behavioral and adaptive technologies that are not reliant on signatures and protecting endpoints at the host level

and on client systems directly addresses the emerging problems that are challenging the enterprise. This is what makes Primary Response unique.

What server applications does Primary Response protect?

Primary Response is application agnostic. Because the SanAPT technology autonomously adapts to any server-based process, Primary Response protects standard applications (such as IIS, Apache and iPlanet Web servers), complex applications (such as Microsoft Exchange, Peoplesoft, SAP, and Oracle) and custom, in-house developed software and applications.

Does Primary Response protect the server as well as the applications running on it?

Yes. Primary Response can profile and protect core operating system services on the server as well as applications running on it. Primary Response includes a set of default applications for securing the operating system with detection protection. Customers can add additional services in the same way they would add new monitored applications.

What types of attacks do Primary Response Server Agents detect and prevent?

Primary Response detects threats that cause server behavior to deviate. This means it will detect a wide range of exploits of program vulnerabilities. Included is a list of some of the major classes of attacks that Primary Response can detect. While this selection is not comprehensive, it is simply intended to highlight the extensive coverage provided by Primary Response:

- Bounds overflows, including buffer overflows, stack overflows, heap overflows, and index overflows
- Code injections, including heap, stack and static memory code injections
- File linkage abuse
- Abuse of incorrectly set permissions
- Abuse of default and sample files
- HTTP header manipulation
- Format string
- Command injection
- Directory traversal
- Abuse of debug functions
- Null bytes
- Off-by-one
- Trojans and backdoors
- Account enumeration
- Race conditions
- Privilege escalation

Because Primary Response is host-based intrusion prevention software, it will not detect abuses that are focused at the network layer, such as connection hijacking and sniffing. Such attacks are best detected or prevented by network-based security solutions such as Network Intrusion Prevention solutions (NIPS).

How does Primary Response adapt to changes in the IT environment such as application changes or system patches?

Primary Response can easily adapt to changing IT environments. First, the security administrator can suspend Primary Response Agents on affected machines, implement the operating system update, application upgrade or security patch, and then resume the agents and instruct them to readapt. The administrator can perform these operations on agents grouped either by application or by machine type (for example, group application patches by application, operating system patches by machine). As the agent readapts, it automatically profiles incremental application changes. While readapting, the agent will not detect, prevent or alert on system call sequence behavior. However, during this time, the agent will continue to detect, prevent and alert on any code injections and buffer overflows, which account for the largest class of attacks in the enterprise. Due to the incremental impact of change management on application profiles, the adaptation to change is typically much quicker than the initial adaptation.

Do I need a unique agent for each server application I want to monitor?

No. Primary Response does not require agents for individual applications. Rather, a single Primary Response Agent can profile multiple applications on a server, including standard applications (such as IIS, Apache and iPlanet Web servers), complex applications (such as Microsoft Exchange, Peoplesoft, SAP and Oracle) and custom, in-house developed software and applications.

We frequently scan and patch our servers. Do we still need Primary Response?

Yes. Vulnerability Scans are good at identifying *known* potential exploits within applications and operating systems across the enterprise. However, Primary Response provides protection from both known and unknown threats, providing a protective shield around servers and applications. Primary Response allows you to schedule security patch deployments and avoid time-consuming security patch fire drills. In addition, Primary Response protects applications you might have that are unsupported and no longer receive patches, for example applications running on Microsoft Windows NT.

Why can't I just install NIDS in front of my servers?

NIDS play an important role in securing network traffic. They are effective at detecting certain types of attacks, such as Denial of Service (DoS) and Distributed Denial of Service (DDoS), and can detect known attacks before they reach servers. However, NIDS do not effectively protect server-based applications even if deployed in front of server farms or individual servers. Because they have no knowledge of application behavior and often no knowledge of what applications are present on which servers, NIDS generate thousands of false alarms forcing administrators to make tradeoffs between accuracy and efficacy by turning off rules and signatures governing their security policies. Moreover, NIDS are incapable of inspecting encrypted traffic such as SSL. Most security experts recommend a layered approach to information security. Primary Response is an application security solution that complements NIDS and other network-layer products, and protects server-based applications more effectively and accurately.

We always penetration test our applications before we deploy them to production servers. Do we still need Primary Response?

Penetration testing will detect application vulnerabilities for known attacks. Primary Response will protect the application from both known and unknown attacks. In addition, Primary Response

provides robust forensic and analytic information that can aid in diagnosing security issues in an application development environment.

How does the Primary Response Management Server communicate with agents?

Primary Response Agents communicate with the Management Server using a web-based authenticated and encrypted protocol (self-certified SSL). Because this is a standards-based protocol, it enables web-based remote management across routers and firewalls. Primary Response Agents do not require new ports to be opened, nor are they listening on existing open ports.

Enforcement

Does Primary Response support Management Groups?

Yes. Primary Response Agents can be managed in user-defined groups. A Primary Response Group is a set of agents and applications protected by the agents. Agents can belong to more than one group, and multiple users can be assigned to manage each group.

Does Primary Response support Role Based User Management?

Yes. Primary Response supports two types of users, Administrators and Group Managers. The Group Manager role provides privileges for managing a group of agents. The Administrator role has full access privileges in Primary Response. Administrators are responsible for global settings, can create groups and group managers and can assign machines, applications and group managers to groups.

How does Primary Response update policies across the enterprise?

Policies are shared across machines and across groups. If a shared policy is updated, all machines using that policy benefit. Groups provide automatic policy updates, pushing the latest settings out to all the machines in the group. For example, if a newly discovered malware program is added to the quarantine list of the malware policy, that malware program will be terminated on all machines using that policy. An administrator can create new policies or replace existing ones by applying policies across management groups.

Visibility

What is required in order to secure servers and underlying applications with Primary Response?

The Primary Response administrator deploys Primary Response Agents to selected servers and designates the applications to profile. This can be accomplished in minutes from the Primary Response Management Console. Once the agents are deployed, applications have out-of-the box protection from code injection and buffer overflow attacks. The agents automatically profile application behavior in order to detect and prevent additional types of attacks. There are no signature libraries to keep up-to-date.

How long does Primary Response take to profile server application behavior?

The adaptation period is autonomously determined by the Primary Response Agent, depending on the type and complexity of the application it is monitoring. The agent monitors the quantity of new system call code paths, which it observes over time, and lowers its adaptation sensitivity

threshold accordingly. Because server-based applications typically perform repetitive tasks, Primary Response can build a profile of normal behavior in several hours. In the case of applications exhibiting more erratic behavior, the agent may profile the application for a few days before it has completely learned the normal application behavior.

Does Primary Response take automatic action on all alerts or can it be configured to only detect attacks and generate alerts?

Primary Response Agents can be set to either detect or prevent attacks. Because these settings are centrally managed, you can easily switch from detection to prevention at any time across any or all profiled applications or protected machines.

What is the impact of Primary Response on the IT staff?

Primary Response can be installed, configured and start protecting PCs, servers and their applications in less than 20 minutes. This allows more effective utilization of IT and security resources. Customer feedback indicates this is a very different experience from rules-only solutions, which require a much higher the level of tuning effort to ensure the appropriate security posture across the enterprise. Additionally, Primary Response's role-based user management and management groups make large deployments easily scalable.

What are the machine alerts in Primary Response?

Primary Response provides administrators with system visibility into their enterprise infrastructure. Administrators receive machine alerts on activities such as agent startup and shutdown, unexpected agent shutdown, agents going offline, and new agent registrations.

Can administrators perform actions responding to alerts they receive?

Certain types of alerts are actionable. Application discovery and malware alerts provide options for handling executables. With application discovery, you can choose to automatically protect a newly discovered application or ignore it. With malware, you can choose to terminate the executable, accept it as a safe executable, or add it to the quarantine list. These settings are then propagated to other machines through sharing of the same policies.

How does Primary Response handle upgrades?

Upgrades and patches to Primary Response, including new agents, can be installed from the Management Console. The Administrator can download the patch from the Sana Security Customer Support site, and then upload them to the Primary Response Management Server from the Management Console.

What are bootless agents?

Primary Response does not require stopping and restarting PCs or servers to install the agents, unlike other IPS products. This capability is known as bootless agents.

What is the Primary Response Agent performance impact on protected PCs and servers?

Primary Response Agents typically consume less than 5% of CPU utilization, up to 10% during sustained attacks. For a detailed list of system requirements, including processor, RAM and disk space requirements by platform, see http://www.sanasecurity.com/products/pr/sys_req.php

Integration

What platforms do Primary Response Server Agents support?

Primary Response Server Agent supports Windows NT 4.0 Server, Windows 2000 Server, Windows 2003 Server, Solaris 8 Server (32-bit and 64-bit) and Solaris 9 Server Operating Systems. For a detailed list of system requirements, including processor, RAM and disk space requirements by platform, see <http://www.sanasecurity.com/products/ppRequirements.php>.

Does Primary Response integrate with LDAP directories?

Yes. Primary Response is designed to support multiple integration points with LDAP directories.

What type of LDAP integration is available with Primary Response?

Primary Response supports multiple integration points with LDAP directory:

- Primary Response user authentication against LDAP
- Primary Response user role authorization based on LDAP
- Primary Response user membership and authorization to manage a group based LDAP
- Machine membership in a management group based on LDAP

You can choose the capabilities that work best for your integration requirements.

What LDAP directories are supported with Primary Response?

Primary Response supports Active Directory 2003 and Sun ONE Directory Server 5.2. Primary Response is designed to work with any standard LDAP based directory.

Can Primary Response integrate with other management consoles?

Yes. In addition to the central management functionality built into the Management Server, Primary Response supports both SNMPv1 and SNMPv2, allowing integration of alerts and system information into third party system management and security information management solutions.

What does Oracle database support allow me to do?

Primary Response stores alerts, system configuration and other data in a repository. The repository can be either a database internal to Primary Response or an Oracle database. With Oracle, administrators can use Oracle database tools to manage data, create detailed, custom reports using reporting packages such as Crystal Reports and integrate Primary Response data with other applications.